

Query shredding: Efficient relational evaluation of queries over nested multisets (extended version)

James Cheney
University of Edinburgh
jcheney@inf.ed.ac.uk

Sam Lindley
University of Edinburgh
Sam.Lindley@ed.ac.uk

Philip Wadler
University of Edinburgh
wadler@inf.ed.ac.uk

ABSTRACT

Nested relational query languages have been explored extensively, and underlie industrial language-integrated query systems such as Microsoft’s LINQ. However, relational databases do not natively support nested collections in query results. This can lead to major performance problems: if programmers write queries that yield nested results, then such systems typically either fail or generate a large number of queries. We present a new approach to query shredding, which converts a query returning nested data to a fixed number of SQL queries. Our approach, in contrast to prior work, handles multiset semantics, and generates an idiomatic SQL:1999 query directly from a normal form for nested queries. We provide a detailed description of our translation and present experiments showing that it offers comparable or better performance than a recent alternative approach on a range of examples.

1. INTRODUCTION

Databases are one of the most important applications of declarative programming techniques. However, relational databases only support queries against flat tables, while programming languages typically provide complex data structures that allow arbitrary combinations of types including nesting of collections (e.g. sets of sets). Motivated by this so-called *impedance mismatch*, and inspired by insights into language design based on monadic comprehensions [33], database researchers introduced nested relational query languages [26, 4, 5] as a generalisation of flat relational queries to allow nesting collection types inside records or other types. Several recent language designs, such as XQuery [25] and PigLatin [23], have further extended these ideas, and they have been particularly influential on language-integrated querying systems such as Kleisli [35], Microsoft’s LINQ in C# and F# [22, 28, 6], Links [8, 20], and Ferry [11].

This paper considers the problem of translating nested queries over nested data to flat queries over a flat representation of nested data, or *query shredding* for short. Our motivation is to support a free combination of the features of nested relational query languages with those of high-level programming languages, particularly systems such as Links, Ferry, and LINQ. All three of these

systems support queries over nested data structures (e.g. records containing nested sets, multisets/bags, or lists) in principle; however, only Ferry supports them in practice. Links and LINQ currently either reject such queries at run-time or execute them inefficiently in-memory by loading unnecessarily large amounts of data or issuing large numbers of queries (sometimes called *query storms* or *avalanches* [12] or the $N + 1$ *query problem*). To construct nested data structures while avoiding this performance penalty, programmers must currently write flat queries (e.g. loading in a superset of the needed source data) and convert the results to nested data structures. Manually reconstructing nested query results is tricky and hard to maintain; it may also mask optimisation opportunities.

In the Ferry system, Grust et al. [11, 12] have implemented shredding for nested list queries by adapting an XQuery-to-SQL translation called *loop-lifting* [15]. Loop-lifting produces queries that make heavy use of advanced On-Line Analytic Processing (OLAP) features of SQL:1999, such as ROW_NUMBER and DENSE_RANK, and to optimise these queries Ferry relies on a SQL:1999 query optimiser called Pathfinder [14].

Van den Bussche [31] proved expressiveness results showing that it is possible in principle to evaluate nested queries over *sets* via multiple flat queries. To strengthen the result, Van den Bussche’s simulation eschews value invention mechanisms such as SQL:1999’s ROW_NUMBER. The downside, however, is that the flat queries can produce results that are quadratically larger than needed to represent sets and may not preserve bag semantics. In particular, for bag semantics the size of a union of two nested sets can be quadratic in the size of the inputs — we give a concrete example in Appendix A.

Query shredding is related to the well-studied *query unnesting* problem [17, 10]. However, most prior work on unnesting only considers SQL queries that contain subqueries in WHERE clauses, not queries returning nested results; the main exception is Fegaras and Maier’s work on query unnesting in a complex object calculus [10].

In this paper, we introduce a new approach to query shredding for nested multiset queries (a case not handled by prior work). Our work is formulated in terms of the Links language, but should be applicable to other language-integrated query systems, such as Ferry and LINQ, or to other complex-object query languages [10]. Figure 1 illustrates the behaviour of Links, Ferry and our approach.

We decompose the translation from nested to flat queries into a series of simpler translations. We leverage prior work on *normalisation* that translates a higher-order query to a *normal form* in which higher-order features have been eliminated [34, 7, 10]. Our algorithm operates on normal forms. We review normalisation in Section 2, and give a running example of our approach in Section 3.

Sections 4–7 present the remaining phases of our approach, which




Figure 1: (a) Default Links behaviour (flat queries) (b) Ferry (nested list queries) (c) Our approach (nested bag queries)

are new. The *shredding* phase translates a single, nested query to a number of flat queries. These queries are organised in a *shredded package*, which is essentially a type expression whose collection type constructors are annotated with queries. The different queries are linked by *indexes*, that is, keys and foreign keys. The shredding translation is presented in Section 4. Shredding leverages the normalisation phase in that we can define translations on types and terms independently. Section 5 shows how to run shredded queries and stitch the results back together to form nested values.

The *let-insertion* phase implements a flat indexing scheme using a let-binding construct and a row-numbering operation. Let-insertion (described in Section 6) is conceptually straightforward, but provides a vital link to proper SQL by providing an implementation of abstract indexes. The final stage is to translate to SQL (Section 7) by flattening records, translating let-binding to SQL's `WITH`, and translating row-numbering to SQL's `ROW_NUMBER`. This phase also flattens nested records; the (standard) details are included in Appendix E for completeness.

We have implemented and experimentally evaluated our approach (Section 8) in comparison with Ulrich's implementation of loop-lifting in Links [30]. Our approach typically performs as well or better than loop-lifting, and can be significantly faster.

Our contribution over prior work can be summarised as follows. Fegaras and Maier [10] show how to unnest complex object queries including lists, sets, and bags, but target a nonstandard nested relational algebra, whereas we target standard SQL. Van den Bussche's simulation [31] translates nested set queries to several relational queries but was used to prove a theoretical result and was not intended as a practical implementation technique, nor has it been implemented and evaluated. Ferry [11, 12] translates nested list queries to several SQL:1999 queries and then tries to simplify the resulting queries using Pathfinder. Sometimes, however, this produces queries with cross-products inside OLAP operations such as `ROW_NUMBER`, which Pathfinder cannot simplify. In contrast, we delay introducing OLAP operations until the last stage, and our experiments show how this leads to much better performance on some queries. Finally, we handle nested multisets, not sets or lists.

This report is the extended version of a SIGMOD 2014 paper, and contains additional details, proofs of correctness, and comparison with related work in appendices.

2. BACKGROUND

We use metavariables x, y, \dots, f, g for *variables*, and c, d, \dots for *constants* and primitive operations. We also use letters t, t', \dots for *table names*, $\ell, \ell', \ell_i, \dots$ for *record labels* and a, b, \dots for *tags*.

We write $M[x := N]$ for capture-avoiding substitution of N for x in M . We write \vec{x} for a vector x_1, \dots, x_n . Moreover, we extend vector notation pointwise to other constructs, writing, for example, $\langle \vec{\ell} = \vec{M} \rangle$ for $\langle \ell_1 = M_1, \dots, \ell_n = M_n \rangle$.

We write: square brackets $[-]$ for the meta level list constructor; $w :: \vec{v}$ for adding the element w onto the front of the list \vec{v} ; $\vec{v} ++ \vec{w}$ for the result of appending the list \vec{w} onto the end of the list \vec{v} ; and *concat* for the function that concatenates a list of lists. We also

$$\begin{aligned}
\mathcal{N}[\![x]\!]_{\rho} &= \rho(x) \\
\mathcal{N}[\![c(X_1, \dots, X_n)]\!]_{\rho} &= [c](\mathcal{N}[\![X_1]\!]_{\rho}, \dots, \mathcal{N}[\![X_n]\!]_{\rho}) \\
\mathcal{N}[\![\lambda x. M]\!]_{\rho} &= \lambda v. \mathcal{N}[\![M]\!]_{\rho[x \mapsto v]} \\
\mathcal{N}[\![M N]\!]_{\rho} &= \mathcal{N}[\![M]\!]_{\rho}(\mathcal{N}[\![N]\!]_{\rho}) \\
\mathcal{N}[\![\langle \ell_i = M_i \rangle_{i=1}^n]\!]_{\rho} &= \langle \ell_i = \mathcal{N}[\![M_i]\!]_{\rho} \rangle_{i=1}^n \\
\mathcal{N}[\![M.\ell]\!]_{\rho} &= \mathcal{N}[\![M]\!]_{\rho}.\ell \\
\mathcal{N}[\![\text{if } L \text{ then } M \text{ else } N]\!]_{\rho} &= \begin{cases} \mathcal{N}[\![M]\!]_{\rho}, & \text{if } \mathcal{N}[\![L]\!]_{\rho} = \text{true} \\ \mathcal{N}[\![N]\!]_{\rho}, & \text{if } \mathcal{N}[\![L]\!]_{\rho} = \text{false} \end{cases} \\
\mathcal{N}[\![\text{return } M]\!]_{\rho} &= [\mathcal{N}[\![M]\!]_{\rho}] \\
\mathcal{N}[\![\emptyset]\!]_{\rho} &= [] \\
\mathcal{N}[\![M \uplus N]\!]_{\rho} &= \mathcal{N}[\![M]\!]_{\rho} \uplus \mathcal{N}[\![N]\!]_{\rho} \\
\mathcal{N}[\![\text{for } (x \leftarrow M) N]\!]_{\rho} &= \text{concat}[\mathcal{N}[\![N]\!]_{\rho[x \mapsto v]} \mid v \leftarrow \mathcal{N}[\![M]\!]_{\rho}] \\
\mathcal{N}[\![\text{empty } M]\!]_{\rho} &= \begin{cases} \text{true}, & \text{if } \mathcal{N}[\![M]\!]_{\rho} = [] \\ \text{false}, & \text{if } \mathcal{N}[\![M]\!]_{\rho} \neq [] \end{cases} \\
\mathcal{N}[\![\text{table } t]\!]_{\rho} &= [t]
\end{aligned}$$

Figure 2: Semantics of λ_{NRC}

make use of the following functions:

$$\begin{aligned}
\text{init } [x_i]_{i=1}^n &= [x_i]_{i=1}^{n-1} \quad \text{last } [x_i]_{i=1}^n = x_n \\
\text{enum}([v_1, \dots, v_m]) &= [\langle 1, v_1 \rangle, \dots, \langle m, v_m \rangle]
\end{aligned}$$

In the meta language we make extensive use of comprehensions, primarily list comprehensions. For instance, $[v \mid x \leftarrow xs, y \leftarrow ys, p]$, returns a copy of v for each pair $\langle x, y \rangle$ of elements of xs and ys such that the predicate p holds. We write $[v_i]_{i=1}^n$ as shorthand for $[v_i \mid 1 \leq i \leq n]$ and similarly, e.g., $\langle \ell_i = M_i \rangle_{i=1}^n$ for $\langle \ell_1 = M_1, \dots, \ell_n = M_n \rangle$.

2.1 Nested relational calculus

We take the higher-order, nested relational calculus (evaluated over bags) as our starting point. We call this λ_{NRC} ; this is essentially a core language for the query components of Links, Ferry, and LINQ. The types of λ_{NRC} include base types (integers, strings, booleans), record types $\langle \vec{\ell} : \vec{A} \rangle$, bag types $\text{Bag } A$, and function types $A \rightarrow B$.

$$\begin{aligned}
\text{Types} \quad A, B &::= O \mid \langle \vec{\ell} : \vec{A} \rangle \mid \text{Bag } A \mid A \rightarrow B \\
\text{Base types} \quad O &::= \text{Int} \mid \text{Bool} \mid \text{String}
\end{aligned}$$

We say that a type is *nested* if it contains no function types and *flat* if it contains only base and record types.

The terms of λ_{NRC} include λ -abstractions, applications, and the standard terms of nested relational calculus.

$$\begin{aligned}
\text{Terms} \quad M, N &::= x \mid c(\vec{M}) \mid \text{table } t \mid \text{if } M \text{ then } N \text{ else } N' \\
&\mid \lambda x. M \mid M N \mid \langle \vec{\ell} = \vec{M} \rangle \mid M.\ell \mid \text{empty } M \\
&\mid \text{return } M \mid \emptyset \mid M \uplus N \mid \text{for } (x \leftarrow M) N
\end{aligned}$$

We assume that the constants and primitive functions include boolean values with negation and conjunction, and integer values with standard arithmetic operations and equality tests. We assume special labels $\#_1, \#_2, \dots$ and encode tuple types $\langle A_1, \dots, A_n \rangle$ as record types $\langle \#_1 : A_1, \dots, \#_n : A_n \rangle$, and similarly tuple terms

$\langle M_1, \dots, M_n \rangle$ as record terms $\langle \#_1 = M_1, \dots, \#_n = M_n \rangle$. We assume fixed signatures $\Sigma(t)$ and $\Sigma(c)$ for tables and constants. The tables are constrained to have flat relation type $(\text{Bag } \langle \ell_1 : O_1, \dots, \ell_n : O_n \rangle)$, and the constants must be of base type or first order n -ary functions $(\langle O_1, \dots, O_n \rangle \rightarrow O)$.

Most language constructs are standard. The \emptyset expression builds an empty bag, return M constructs a singleton, and $M \uplus N$ builds the bag union of two collections. The for $(x \leftarrow M) N$ comprehension construct iterates over a bag obtained by evaluating M , binds x to each element, evaluates N to another bag for each such binding, and takes the union of the results. The expression empty M evaluates to true if M evaluates to an empty bag, false otherwise.

λ_{NRC} employs a standard type system similar to that presented in other work [35, 20, 6]. We will also employ several typed intermediate languages and translations mapping λ_{NRC} to SQL. All of these (straightforward) type systems are omitted due to space limits; they will be available in the full version of this paper.

Semantics. We give a denotational semantics in terms of lists. Though we wish to preserve bag semantics, we interpret object-level bags as meta-level lists. For meta-level values v and v' , we consider v and v' equivalent as multisets if they are equal up to permutation of list elements. We use lists mainly so that we can talk about order-sensitive operations such as `row_number`.

We interpret base types as integers, booleans and strings, function types as functions, record types as records, and bag types as lists. For each table $t \in \text{dom}(\Sigma)$, we assume a fixed interpretation $\llbracket t \rrbracket$ of t as a list of records of type $\Sigma(t)$. In SQL, tables do not have a list semantics by default, but we can impose one by choosing a canonical ordering for the rows of the table. We order by all of the columns arranged in lexicographic order (assuming linear orderings on field names and base types).

We assume fixed interpretations $\llbracket c \rrbracket$ for the constants and primitive operations. The semantics of nested relational calculus are shown in Figure 2. The (standard) typing rules are given in Appendix B. We let ρ range over environments mapping variables to values, writing ε for the empty environment and $\rho[x \mapsto v]$ for the extension of ρ with x bound to v .

2.2 Query normalisation

In Links, query normalisation is an important part of the execution model [7, 20]. Links currently supports only queries mapping flat tables to flat results, or *flat-flat* queries. When a subexpression denoting a query is evaluated, the subexpression is first normalised and then converted to SQL, which is sent to the database for evaluation; the tuples received in response are then converted into a Links value and normal execution resumes (see Figure 1(a)).

For *flat-nested* queries that read from flat tables and produce a nested result value, our normalisation procedure is similar to the one currently used in Links [20], but we hoist all conditionals into the nearest enclosing comprehension as where clauses. This is a minor change; the modified algorithm is given in the full version of this paper. The resulting normal forms are:

Query terms	$L ::= \uplus \vec{C}$
Comprehensions	$C ::= \text{for } (\vec{G} \text{ where } X) \text{ return } M$
Generators	$G ::= x \leftarrow t$
Normalised terms	$M, N ::= \overline{X} \mid R \mid L$
Record terms	$R ::= \langle \ell = \vec{M} \rangle$
Base terms	$X ::= x.\ell \mid c(\vec{X}) \mid \text{empty } L$

Any closed flat-nested query can be converted to an equivalent term in the above normal form.

$\llbracket \text{departments} \rrbracket =$		$\llbracket \text{employees} \rrbracket =$			
(id)	name	(id)	dept	name	salary
1	Product	1	Product	Alex	20000
2	Quality	2	Product	Bert	900
3	Research	3	Research	Cora	50000
4	Sales	4	Research	Drew	60000
		5	Sales	Erik	2000000
		6	Sales	Fred	700
		7	Sales	Gina	100000

$\llbracket \text{tasks} \rrbracket =$			$\llbracket \text{contacts} \rrbracket =$			
(id)	employee	task	(id)	dept	name	client
1	Alex	build	1	Product	Pam	false
2	Bert	build	2	Product	Pat	true
3	Cora	abstract	3	Research	Rob	false
4	Cora	build	4	Research	Roy	false
5	Cora	call	5	Sales	Sam	false
6	Cora	dissemble	6	Sales	Sid	false
7	Cora	enthuse	7	Sales	Sue	true
8	Drew	abstract				
9	Drew	enthuse				
10	Erik	call				
11	Erik	enthuse				
12	Fred	call				
13	Gina	call				
14	Gina	dissemble				

Figure 3: Sample data

THEOREM 1. *Given a closed flat-nested query $\vdash M : \text{Bag } A$, there exists a normalisation function $\text{norm}_{\text{Bag } A}$, mapping each M to an equivalent normal form $\text{norm}_{\text{Bag } A}(M)$.*

The normalisation algorithm and correctness proof are similar to those in previous papers [7, 20, 6]. The details are given in Appendix C. The normal forms above can also be viewed as an SQL-like language allowing relation-valued attributes (similar to the complex-object calculus of Fegaras and Maier [10]). Thus, our results can also be used to support nested query results in an SQL-like query language. In this paper, however, we focus on the functional core language based on comprehensions, as illustrated in the next section.

3. RUNNING EXAMPLE

To motivate and illustrate our work, we present an extended example showing how our shredding translation could be used to provide useful functionality to programmers working in LINQ using F#, Ferry or Links. We first describe the code the programmer would actually write and the results the system produces. Throughout the rest of the paper, we return to this example to illustrate how the shredding translation works.

Consider a flat database schema Σ for an organisation:

```
tasks(employee : String, task : String)
employees(dept : String, name : String, salary : Int)
contacts(dept : String, name : String, client : Bool)
departments(name : String)
```

Each department has a name, a collection of employees, and a collection of external contacts. Each employee has a name, salary and a collection of tasks. Some contacts are clients. Figure 3 shows a small instance of this schema. For convenience, we also assume every table has an integer-valued key *id*.

In λ_{NRC} , queries of the form for \dots where \dots return \dots are natively supported. These are *comprehensions* as found in XQuery or functional programming languages, and they generalise idiomatic

SQL SELECT FROM WHERE queries [4]. Unlike SQL, we can use (nonrecursive) functions to define queries with parameters, or parts of queries, and freely combine them. For example, the following functions define useful query patterns over the above schema:

```
tasksOfEmp e = for (t ← tasks)
  where (t.employee = e.name)
  return t.task

contactsOfDept d = for (c ← contacts)
  where (d.dept = c.dept)
  return {name = c.name, client = c.client}

employeesByTask t =
  for (e ← employees, d ← departments)
  where (e.name = t.employee ∧ e.dept = d.dept)
  return {b = e.employee, c = d.dept}
```

Nested queries allow free mixing of collection (bag) types with record or base types. For example, the following query

```
employeesOfDept d = for (e ← employees)
  where (d.dept = e.dept)
  return {name = e.name, salary = e.salary,
         tasks = tasksOfEmp e}
```

returns a nested result: a collection of employees in a department, each with an associated collection of tasks. That is, its return type is $\text{Bag} \langle \text{name} : \text{String}, \text{salary} : \text{Int}, \text{tasks} : \text{Bag String} \rangle$

Consider the following nested schema for organisations:

```
Task = String
Employee = {name : String, salary : Int, tasks : Bag Task}
Contact = {name : String, client : Bool}
Department = {name : String, employees : Bag Employee,
              contacts : Bag Contact}
Organisation = Bag Department
```

Using some of the above functions we can write a query Q_{org} that maps data in the flat schema Σ to the nested type *Organisation*, as follows:

```
 $Q_{org} = \text{for } (d \leftarrow \text{departments})$ 
  return ( $\langle \text{name} = d.\text{name},$ 
          $\text{employees} = \text{employeesOfDept } d,$ 
          $\text{contacts} = \text{contactsOfDept } d \rangle$ )
```

We can also define and use higher-order functions to build queries, such as the following:

```
filter p xs = for (x ← xs) where (p(x)) return x
any xs p = ¬(empty(for (x ← xs) where (p(x)) return {}))
all xs p = ¬(any xs (λx.¬(p(x))))
contains xs u = any xs (λx.x = u)
```

To illustrate the main technical challenges of shredding, we consider a query with two levels of nesting and a union operation.

Suppose we wish to find for each department a collection of people of interest, both employees and contacts, along with a list of the tasks they perform. Specifically, we are interested in those employees that earn less than a thousand euros and those who earn more than a million euros, call them *outliers*, along with those contacts who are clients. The following code defines poor, rich, outliers, and clients:

```
isPoor x = x.salary < 1000
isRich x = x.salary > 1000000
outliers xs = filter (λx.isRich x ∨ isPoor x) xs
clients xs = filter (λx.x.client) xs
```

We also introduce a convenient higher-order function that uses its f parameter to initialise the tasks of its elements:

```
getTasks xs f = for (x ← xs) return {name = x.name, tasks = f x}
```

Using the above operations, the query Q returns each department, the outliers and clients associated with that department, and their tasks. We assign the special task “buy” to clients.

```
 $Q(\text{organisation}) =$ 
  for ( $x \leftarrow \text{organisation}$ )
  return ( $\langle \text{department} = x.\text{name},$ 
          $\text{people} =$ 
            $\text{getTasks}(\text{outliers}(x.\text{employees})) (\lambda y.y.\text{tasks})$ 
            $\sqcup \text{getTasks}(\text{clients}(x.\text{contacts})) (\lambda y.\text{return "buy"}) \rangle$ )
```

The result type of Q is:

```
Result = Bag {department : String,
              people : Bag {name : String, tasks : Bag String}}
```

We can compose Q with Q_{org} to form a query $Q(Q_{org})$ from the flat data stored in Σ to the nested *Result*. The normal form of this composed query, which we call Q_{comp} , is as follows:

```
 $Q_{comp} = \text{for } (x \leftarrow \text{departments})$ 
  return (
     $\langle \text{department} = x.\text{name},$ 
     $\text{people} =$ 
      ( $\text{for } (y \leftarrow \text{employees}) \text{ where } (x.\text{name} = y.\text{dept} \wedge$ 
        $(y.\text{salary} < 1000 \vee y.\text{salary} > 1000000))$ 
       return ( $\langle \text{name} = y.\text{name},$ 
               $\text{tasks} = \text{for } (z \leftarrow \text{tasks})$ 
                where ( $z.\text{employee} = y.\text{name}$ )
                return  $z.\text{task}$ 
               $\rangle$ )
        $\sqcup (\text{for } (y \leftarrow \text{contacts})$ 
          where ( $x.\text{name} = y.\text{dept} \wedge y.\text{client}$ )
          return ( $\langle \text{name} = y.\text{name},$ 
                  $\text{tasks} = \text{return "buy"}$ 
                  $\rangle$ 
               $\rangle$ 
      )
    )
```

The result of running Q_{comp} on the data in Figure 3 is:

```
{[department = "Product",
  people = [{name = "Bert", tasks = ["build"]},
            {name = "Pat", tasks = ["buy"]}]}]
[department = "Research", people = {}],
[department = "Quality", people = {}],
[department = "Sales",
  people = [{name = "Erik", tasks = ["call", "enthuse"]},
            {name = "Fred", tasks = ["call"]},
            {name = "Sue", tasks = ["buy"]}]}]
```

Now, however, we are faced with a problem: SQL databases do not directly support nested multisets (or sets). Our shredding translation, like Van den Bussche’s simulation for sets [31] and Grust et al.’s for lists [12], can translate a normalised query such as $Q_{comp} : \text{Result}$ that maps flat input Σ to nested output *Result* to a fixed number of flat queries $q_1 : \text{Result}_1, \dots, q_n : \text{Result}_n$ whose results can be combined via a *stitching* operation $Q_{stitch} : \text{Result}_1 \times \dots \times \text{Result}_n \rightarrow \text{Result}$. Thus, we can simulate the query Q_{comp} by running q_1, \dots, q_n remotely on the database and stitching the results together using Q_{stitch} . The number of intermediate queries n is the *nesting degree* of *Result*, that is, the number of collection type constructors in the result type. For example, the nesting degree of $\text{Bag} \langle A : \text{Bag Int}, B : \text{Bag String} \rangle$ is 3. The nesting degree of the type *Result* is also 3, which means Q can be shredded into three flat queries.

The basic idea is straightforward. Whenever a nested bag appears in the output of a query, we generate an index that uniquely identifies the current context. Then a separate query produces the contents of the nested bag, where each element is paired up with its parent index. Each inner level of nesting requires a further query.

We will illustrate by showing the results of the three queries and how they can be combined to reconstitute the desired nested result.

The outer query q_1 contains one entry for each department, with an index $\langle a, id \rangle$ in place of each nested collection:

$$r_1 = [\langle \text{department} = \text{"Product"}, \text{people} = \langle a, 1 \rangle \rangle, \\ \langle \text{department} = \text{"Quality"}, \text{people} = \langle a, 2 \rangle \rangle, \\ \langle \text{department} = \text{"Research"}, \text{people} = \langle a, 3 \rangle \rangle, \\ \langle \text{department} = \text{"Sales"}, \text{people} = \langle a, 4 \rangle \rangle]$$

The second query q_2 generates the data needed for the **people** collections:

$$r_2 = [\langle \langle a, 1 \rangle, \langle \text{name} = \text{"Bert"}, \text{tasks} = \langle b, 1, 2 \rangle \rangle, \\ \langle \langle a, 4 \rangle, \langle \text{name} = \text{"Erik"}, \text{tasks} = \langle b, 4, 5 \rangle \rangle, \\ \langle \langle a, 4 \rangle, \langle \text{name} = \text{"Fred"}, \text{tasks} = \langle b, 4, 6 \rangle \rangle, \\ \langle \langle a, 1 \rangle, \langle \text{name} = \text{"Pat"}, \text{tasks} = \langle d, 1, 2 \rangle \rangle, \\ \langle \langle a, 4 \rangle, \langle \text{name} = \text{"Sue"}, \text{tasks} = \langle d, 4, 7 \rangle \rangle]$$

The idea is to ensure that we can stitch the results of q_1 together with the results of q_2 by joining the *inner indexes* of q_1 (bound to the **people** field of each result) with the *outer indexes* of q_2 (bound to the first component of each result). In both cases the static components of these indexes are the same tag a . Joining the **people** field of q_1 to the outer index of q_2 correctly associates each person with the appropriate department.

Finally, let us consider the results of the innermost query q_3 for generating the bag bound to the **tasks** field:

$$r_3 = [\langle \langle b, 1, 2 \rangle, \text{"build"} \rangle, \langle \langle b, 4, 5 \rangle, \text{"call"} \rangle, \langle \langle b, 4, 5 \rangle, \text{"enthuse"} \rangle, \\ \langle \langle b, 4, 6 \rangle, \text{"call"} \rangle, \langle \langle d, 1, 2 \rangle, \text{"buy"} \rangle, \langle \langle d, 4, 7 \rangle, \text{"buy"} \rangle]$$

Recall that q_2 returns further inner indexes for the tasks associated with each person. The two halves of the union have different static indexes for the tasks b and d , because they arise from different comprehensions in the source term. Furthermore, the dynamic index now consists of two **id** fields ($x.\text{id}$ and $y.\text{id}$) in each half of the union. Thus, joining the **tasks** field of q_2 to the outer index of q_3 correctly associates each task with the appropriate outlier.

Note that each of the queries q_1, q_2, q_3 produces records that contain other records as fields. This is not strictly allowed by SQL, but it is straightforward to simulate such nested records by rewriting to a query with no nested collections in its result type; this is similar to Van den Bussche’s simulation [31]. However, this approach incurs extra storage and query-processing cost. Later in the paper, we explore an alternative approach which collapses the indexes at each level to a pair $\langle a, i \rangle$ of static index and a single “surrogate” integer, similarly to Ferry’s approach [12]. For example, using this approach we could represent the results of q_2 and q_3 as follows:

$$r'_2 = [\langle \langle a, 1 \rangle, \langle \text{name} = \text{"Bert"}, \text{tasks} = \langle b, 1 \rangle \rangle, \\ \langle \langle a, 4 \rangle, \langle \text{name} = \text{"Erik"}, \text{tasks} = \langle b, 2 \rangle \rangle, \\ \langle \langle a, 4 \rangle, \langle \text{name} = \text{"Fred"}, \text{tasks} = \langle b, 3 \rangle \rangle, \\ \langle \langle a, 1 \rangle, \langle \text{name} = \text{"Pat"}, \text{tasks} = \langle d, 1 \rangle \rangle, \\ \langle \langle a, 4 \rangle, \langle \text{name} = \text{"Sue"}, \text{tasks} = \langle d, 2 \rangle \rangle], \\ r'_3 = [\langle \langle b, 1 \rangle, \text{"build"} \rangle, \langle \langle b, 2 \rangle, \text{"call"} \rangle, \langle \langle b, 2 \rangle, \text{"enthuse"} \rangle, \\ \langle \langle b, 3 \rangle, \text{"call"} \rangle, \langle \langle d, 1 \rangle, \text{"buy"} \rangle, \langle \langle d, 2 \rangle, \text{"buy"} \rangle]$$

The rest of this paper gives the details of the shredding translation, explains how to stitch the results of shredded queries back together, and shows how to use `row_number` to avoid the space overhead of indexes. We will return to the above example throughout the paper, and we will use Q_{orig}, Q and other queries based on this example in the experimental evaluation.

4. SHREDDING TRANSLATION

As a pre-processing step, we annotate each comprehension body in a normalised term with a unique name a — the static component of an index. We write the annotations as superscripts, for example:

$$\text{for } (\vec{G} \text{ where } X) \text{ return}^a M$$

In order to shred nested queries, we introduce an abstract type *Index of indexes* for maintaining the correspondence between outer and inner queries. An index $a \diamond d$ has a static component a and a dynamic component d . The static component a links the index to the corresponding `returna` in the query. The dynamic component identifies the current bindings of the variables in the comprehension.

Next, we modify types so that bag types have an explicit index component and we use indexes to replace nested occurrences of bags within other bags:

$$\begin{array}{ll} \text{Shredded types} & A, B ::= \text{Bag } \langle \text{Index}, F \rangle \\ \text{Flat types} & F ::= O \mid \langle \ell : \vec{F} \rangle \mid \text{Index} \end{array}$$

We also adapt the syntax of terms to incorporate indexes. After shredding, terms will have the following forms:

$$\begin{array}{ll} \text{Query terms} & L, M ::= \biguplus \vec{C} \\ \text{Comprehensions} & C ::= \text{return}^a \langle I, N \rangle \\ & \quad \mid \text{for } (\vec{G} \text{ where } X) C \\ \text{Generators} & G ::= x \leftarrow t \\ \text{Inner terms} & N ::= X \mid R \mid I \\ \text{Record terms} & R ::= \langle \ell = \vec{N} \rangle \\ \text{Base terms} & X ::= x.\ell \mid c(\vec{X}) \mid \text{empty } L \\ \text{Indexes} & I, J ::= a \diamond d \\ \text{Dynamic indexes} & d ::= \text{out} \mid \text{in} \end{array}$$

A comprehension is now constructed from a sequence of generator clauses of the form `for` (\vec{G} where X) followed by a body of the form `returna` $\langle I, N \rangle$. Each level of nesting gives rise to such a generator clause. The body always returns a pair $\langle I, N \rangle$ of an outer index I , denoting where the result values from the shredded query should be spliced into the final nested result, and a (flat) inner term N . Records are restricted to contain inner terms, which are either base types, records, or indexes, which replace nested multisets. We assume a distinguished top level static index \top , which allows us to treat all levels uniformly. Each shredded term is associated with an outer index `out` and an inner index `in`. (In fact `out` only appears in the left component of a comprehension body, and `in` only appears in the right component of a comprehension body. These properties will become apparent when we specify the shredding transformation on terms.)

4.1 Shredding types and terms

We use *paths* to point to parts of types.

$$\text{Paths } p ::= \epsilon \mid \downarrow.p \mid \ell.p$$

The empty path is written ϵ . A path p can be extended by traversing a bag constructor ($\downarrow.p$) or selecting a label ($\ell.p$). We will sometimes write $p.\downarrow$ for the path p with \downarrow appended at the end and similarly for $p.\ell$; likewise, we will write $p.\vec{\ell}$ for the path p with all the labels of $\vec{\ell}$ appended. The function $\text{paths}(A)$ defines the set of paths to bag types in a type A :

$$\begin{aligned} \text{paths}(O) &= \{\} \\ \text{paths}(\langle \ell_i : A_i \rangle_{i=1}^n) &= \bigcup_{i=1}^n \{\ell_i.p \mid p \leftarrow \text{paths}(A_i)\} \\ \text{paths}(\text{Bag } A) &= \{\epsilon\} \cup \{\downarrow.p \mid p \leftarrow \text{paths}(A)\} \end{aligned}$$

We now define a shredding translation on types. This is defined in terms of the *inner shredding* $\llbracket A \rrbracket$, a flat type that represents the contents of a bag.

$$\begin{aligned} \llbracket O \rrbracket &= O \\ \llbracket \langle \ell_i : A_i \rangle_{i=1}^n \rrbracket &= \langle \ell_i : \llbracket A_i \rrbracket \rangle_{i=1}^n \\ \llbracket \text{Bag } A \rrbracket &= \text{Index} \end{aligned}$$

$$\begin{aligned}
\llbracket L \rrbracket_p &= \mathcal{M}(\llbracket L \rrbracket_{\top, p}^*) \\
\llbracket \biguplus_{i=1}^n C_i \rrbracket_{a, p}^* &= \text{concat}(\llbracket C_i \rrbracket_{a, p}^* \rrbracket_{i=1}^n) \\
\llbracket \langle \ell_i = M_i \rangle_{i=1}^n \rrbracket_{a, \ell_j, p}^* &= \llbracket M_j \rrbracket_{a, p}^* \\
\llbracket \text{for } (\vec{G} \text{ where } X) \text{ return}^b M \rrbracket_{a, \epsilon}^* &= [\text{for } (\vec{G} \text{ where } X) \text{ return}^b \langle a \diamond \text{out}, \llbracket M \rrbracket_b \rangle] \\
\llbracket \text{for } (\vec{G} \text{ where } X) \text{ return}^b M \rrbracket_{a, \downarrow, p}^* &= [\text{for } (\vec{G} \text{ where } X) C \mid C \leftarrow \llbracket M \rrbracket_{b, p}^*]
\end{aligned}$$

$$\begin{aligned}
\llbracket x.\ell \rrbracket_a &= x.\ell \\
\llbracket c(\llbracket X_i \rrbracket_{i=1}^n) \rrbracket_a &= c(\llbracket \llbracket X_i \rrbracket_a \rrbracket_{i=1}^n) \\
\llbracket \text{empty } L \rrbracket_a &= \text{empty } \llbracket L \rrbracket_\epsilon \\
\llbracket \langle \ell_i = M_i \rangle_{i=1}^n \rrbracket_a &= \langle \ell_i = \llbracket M_i \rrbracket_a \rrbracket_{i=1}^n \\
\llbracket L \rrbracket_a &= a \diamond \text{in}
\end{aligned}$$

Figure 4: Shredding translation on terms

Given a path $p \in \text{paths}(A)$, the type $\llbracket A \rrbracket_p$ is the *outer shredding* of A at p . It corresponds to the bag at path p in A .

$$\begin{aligned}
\llbracket \text{Bag } A \rrbracket_\epsilon &= \text{Bag } \langle \text{Index}, \llbracket A \rrbracket \rangle \\
\llbracket \text{Bag } A \rrbracket_{\downarrow, p} &= \llbracket A \rrbracket_p \\
\llbracket \langle \ell : \hat{A} \rangle \rrbracket_{\ell_i, p} &= \llbracket A \rrbracket_p
\end{aligned}$$

For example, consider the result type *Result* from Section 3. Its nesting degree is 3, and its paths are:

$$\text{paths}(\text{Result}) = \{\epsilon, \downarrow.\text{people}.\epsilon, \downarrow.\text{people}.\downarrow.\text{tasks}.\epsilon\}$$

We can shred *Result* in three ways using these three paths, yielding three shredded types:

$$\begin{aligned}
A_1 &= \llbracket \text{Result} \rrbracket_\epsilon \\
A_2 &= \llbracket \text{Result} \rrbracket_{\downarrow.\text{people}.\epsilon} \\
A_3 &= \llbracket \text{Result} \rrbracket_{\downarrow.\text{people}.\epsilon.\downarrow.\text{tasks}.\epsilon}
\end{aligned}$$

or equivalently:

$$\begin{aligned}
A_1 &= \text{Bag } \langle \text{Index}, \langle \text{department} : \text{String}, \text{people} : \text{Index} \rangle \rangle \\
A_2 &= \text{Bag } \langle \text{Index}, \langle \text{name} : \text{String}, \text{tasks} : \text{Index} \rangle \rangle \\
A_3 &= \text{Bag } \langle \text{Index}, \text{String} \rangle
\end{aligned}$$

The shredding translation on terms $\llbracket L \rrbracket_p$ is given in Figure 4. This takes a term L and a path p and gives a query $\llbracket L \rrbracket_p$ that computes a result of type $\llbracket A \rrbracket_p$, where A is the type of L . The auxiliary translation $\llbracket M \rrbracket_{a, p}^*$ returns the shredded comprehensions of M along path p with outer static index a . The auxiliary translation $\llbracket M \rrbracket_a$ produces a flat representation of M with inner static index a . Note that the shredding translation is linear in time and space. Observe that for emptiness tests we need only the top-level query.

Continuing the example, we can shred Q_{comp} in three ways, yielding shredded queries:

$$\begin{aligned}
q_1 &= \llbracket Q_{\text{comp}} \rrbracket_\epsilon \\
q_2 &= \llbracket Q_{\text{comp}} \rrbracket_{\downarrow.\text{people}.\epsilon} \\
q_3 &= \llbracket Q_{\text{comp}} \rrbracket_{\downarrow.\text{people}.\epsilon.\downarrow.\text{tasks}.\epsilon}
\end{aligned}$$

or equivalently:

$$\begin{aligned}
q_1 &= \text{for } (x \leftarrow \text{departments}) \\
&\quad \text{return}^a \langle \top \diamond 1, \langle \text{department} = x.\text{name}, \text{people} = a \diamond \text{in} \rangle \rangle \\
q_2 &= (\text{for } (x \leftarrow \text{departments}) \\
&\quad \text{for } (y \leftarrow \text{employees}) \text{ where } (x.\text{name} = y.\text{dept} \wedge \\
&\quad \quad (y.\text{salary} < 1000 \vee y.\text{salary} > 1000000)) \\
&\quad \text{return}^b \langle \langle a \diamond \text{out}, \langle \text{name} = y.\text{name}, \text{tasks} = b \diamond \text{in} \rangle \rangle \rangle) \\
&\quad \uplus (\text{for } (x \leftarrow \text{departments}) \\
&\quad \text{for } (y \leftarrow \text{contacts}) \text{ where } (x.\text{name} = y.\text{dept} \wedge y.\text{client}) \\
&\quad \text{return}^d \langle \langle a \diamond \text{out}, \langle \text{name} = y.\text{name}, \text{tasks} = d \diamond \text{in} \rangle \rangle \rangle) \\
q_3 &= (\text{for } (x \leftarrow \text{departments}) \\
&\quad \text{for } (y \leftarrow \text{employees}) \text{ where } (x.\text{name} = y.\text{dept} \wedge \\
&\quad \quad (y.\text{salary} < 1000 \vee y.\text{salary} > 1000000)) \\
&\quad \text{for } (z \leftarrow \text{tasks}) \text{ where } (z.\text{employee} = y.\text{employee}) \\
&\quad \text{return}^c \langle \langle b \diamond \text{out}, z.\text{task} \rangle \rangle) \\
&\quad \uplus (\text{for } (x \leftarrow \text{departments}) \\
&\quad \text{for } (y \leftarrow \text{contacts}) \text{ where } (x.\text{name} = y.\text{dept} \wedge y.\text{client}) \\
&\quad \text{return}^e \langle \langle d \diamond \text{out}, \text{"buy"} \rangle \rangle)
\end{aligned}$$

As a sanity check, we show that well-formed normalised terms shred to well-formed shredded terms of the appropriate shredded types. We discuss other correctness properties of the shredding translation in Section 5. Typing rules for shredded terms are shown in Appendix B.

THEOREM 2. *Suppose L is a normalised flat-nested query with $\vdash L : A$ and $p \in \text{paths}(A)$, then $\vdash \llbracket L \rrbracket_p : \llbracket A \rrbracket_p$.*

4.2 Shredded packages

To maintain the relationship between shredded terms and the structure of the nested result they are meant to construct, we use *shredded packages*. A shredded package \hat{A} is a nested type with annotations, denoted $(-)^a$, attached to each bag constructor.

$$\hat{A} ::= O \mid \langle \ell : \hat{A} \rangle \mid (\text{Bag } \hat{A})^a$$

For a given package, the annotations are drawn from the same set. We write $\hat{A}(S)$ to denote a shredded package with annotations drawn from the set S . We sometimes omit the type parameter when it is clear from context. For shredded packages are shown in Appendix B.

Given a shredded package \hat{A} , we can erase its annotations to obtain its underlying type.

$$\begin{aligned}
\text{erase}(O) &= O \\
\text{erase}(\langle \ell_i : \hat{A}_i \rangle_{i=1}^n) &= \langle \ell_i : \text{erase}(\hat{A}_i) \rangle_{i=1}^n \\
\text{erase}((\text{Bag } \hat{A})^a) &= \text{Bag } (\text{erase}(\hat{A}))
\end{aligned}$$

Given a type A and a shredding function $f : \text{paths}(A) \rightarrow S$, we can construct a shredded package $\hat{A}(S)$.

$$\begin{aligned}
\text{package}_f(A) &= \text{package}_{f, \epsilon}(A) \\
\text{package}_{f, p}(O) &= O \\
\text{package}_{f, p}(\langle \ell_i : \hat{A}_i \rangle_{i=1}^n) &= \langle \ell_i : \text{package}_{f, p, \ell_i}(\hat{A}_i) \rangle_{i=1}^n \\
\text{package}_{f, p}(\text{Bag } A) &= (\text{Bag } (\text{package}_{f, p, \downarrow}(A)))^{f(p)}
\end{aligned}$$

Using *package*, we lift the type- and term-level shredding functions $\llbracket - \rrbracket$ to produce shredded packages, where each annotation contains the shredded version of the input type or query along the path to the associated bag constructor.

$$\begin{aligned}
\text{shred}_B(A) &= \text{package}_{(\llbracket B \rrbracket_-)}(A) \\
\text{shred}_L(A) &= \text{package}_{(\llbracket L \rrbracket_-)}(A)
\end{aligned}$$

For example, the shredded package for the *Result* type from Section 3 is:

$$\begin{aligned}
\text{Shred}_{\text{Result}}(\text{Result}) &= \\
&\text{Bag } (\text{department} : \text{String}, \\
&\quad \text{people} : \text{Bag } (\text{name} : \text{String}, \\
&\quad \quad \text{tasks} : (\text{Bag } \text{String})^{A_3})^{A_2})^{A_1}
\end{aligned}$$

where A_1, A_2 , and A_3 are as shown in Section 4.1. Shredding the normalised query Q' gives the same package, except the type annotations A_1, A_2, A_3 become queries q_1, q_2, q_3 .

Again, as a sanity check we show that erasure is the left inverse of type shredding and that term-level shredding operations preserve types.

THEOREM 3. *For any type A , we have $\text{erase}(\text{shred}_A(A)) = A$. Furthermore, if L is a closed, normalised, flat-nested query such that $\vdash L : A$ then $\vdash \text{shred}_L(A) : \text{shred}_A(A)$.*

5. QUERY EVALUATION AND STITCHING

Having shredded a normalised nested query, we can then run all of the resulting shredded queries separately. If we stitch the shredded results together to form a nested result, then we obtain the same nested result as we would obtain by running the nested query directly. In this section we describe how to run shredded queries and stitch their results back together to form a nested result.

5.1 Evaluating shredded queries

The semantics of shredded queries $S[-]$ is given in Figure 5. Apart from the handling of indexes, it is much like the semantics for nested queries given in Figure 2. To allow different implementations of indexes, we first define a canonical notion of index, and then parameterise the semantics by the concrete type of indexes X and a function $\text{index} : \text{Index} \rightarrow X$ mapping each canonical index to a distinct concrete index. A canonical index $a \diamond \iota$ comprises static index a and dynamic index ι , where the latter is a list of positive natural numbers. For now we take concrete indexes simply to be canonical indexes, and index to be the identity function. We consider other definitions of index in Section 6.

The current dynamic index ι is threaded through the semantics in tandem with the environment ρ . The former encodes the position of each of the generators in the current comprehension and allows us to invoke index to construct a concrete index. The outer index at dynamic index $\iota.i$ is ι ; the inner index is $\iota.i$. In order for a comprehension to generate dynamic indexes we use the function enum (introduced in Section 2) that takes a list of elements and returns the same list with the element number paired up with each source element.

Running a shredded query yields a list of pairs of indexes and shredded values.

$$\begin{array}{ll} \text{Results} & s ::= [\langle I_1, w_1 \rangle, \dots, \langle I_m, w_m \rangle] \\ \text{Flat values} & w ::= c \mid \langle \ell_1 = w_1, \dots, \ell_n = w_n \rangle \mid I \end{array}$$

Given a shredded package $\hat{A}(S)$ and a function $f : S \rightarrow T$, we can map f over the annotations to obtain a new shredded package $\hat{A}'(T)$ such that $\text{erase}(\hat{A}) = \text{erase}(\hat{A}')$.

$$\begin{aligned} \text{pmap}_f(O) &= O \\ \text{pmap}_f(\langle \ell_i : \hat{A}_i \rangle_{i=1}^n) &= \langle \ell_i : \text{pmap}_f(\hat{A}_i) \rangle_{i=1}^n \\ \text{pmap}_f((\text{Bag } \hat{A})^\alpha) &= ((\text{Bag } \text{pmap}_f(\hat{A})))^{\text{f}(\alpha)} \end{aligned}$$

The semantics of a shredded query package is a *shredded value package* containing indexed results for each shredded query. For each type A we define $\mathcal{H}[A] = \text{shred}_A(A)$ and for each flat-nested, closed $\vdash L : A$ we define $\mathcal{H}[L]_A : \mathcal{H}[A]$ as $\text{pmap}_{S[-]}(\text{shred}_L(A))$. In other words, we first construct the shredded query package $\text{shred}_L(A)$, then apply the shredded semantics $S[q]$ to each query q in the package.

For example, here is the shredded package that we obtain after running the normalised query Q_{comp} from Section 2.2:

$$\mathcal{H}[Q_{\text{comp}}]_A = \text{Bag } \langle \text{department} : \text{String}, \\ \text{people} : \text{Bag } \langle \text{name} : \text{String}, \\ \text{tasks} : (\text{Bag } \text{String})^{r_3} \rangle^{r_2} \rangle^{r_1}$$

where r_1 , r_2 , and r_3 are as in Section 3 except that indexes are of the form $a \diamond 1.2.3$ instead of $\langle a, 1, 2, 3 \rangle$.

5.2 Stitching shredded query results together

A shredded value package can be *stitched* back together into a nested value, preserving annotations, as follows:

$$\begin{aligned} \text{stitch}(\hat{A}) &= \text{stitch}_{\top \diamond 1}(\hat{A}) \\ \text{stitch}_c(O) &= c \\ \text{stitch}_r(\langle \ell_i : \hat{A}_i \rangle_{i=1}^n) &= \langle \ell_i = \text{stitch}_{r.\ell_i}(\hat{A}_i) \rangle_{i=1}^n \\ \text{stitch}_I((\text{Bag } \hat{A})^s) &= [(\text{stitch}_w(\hat{A})) \mid \langle I, w \rangle \leftarrow s] \end{aligned}$$

The flat value parameter w to the auxiliary function $\text{stitch}_w(-)$ specifies which values to stitch along the current path.

Resuming our running example, once the results $r_1 : A_1, r_2 : A_2, r_3 : A_3$ have been evaluated on the database, they are shipped back to the host system where we can run the following code in-memory to stitch the three tables together into a single value: the result of the original nested query. The code for this query Q_{stitch} follows the same idea as the query Q_{org} that constructs the nested *organisation* from Σ .

```
for (x ← r1)
  return ((department = x.name,
    people = for ((i, y) ← r2)
      where (x.people = i))
    return ((name = y.name,
      tasks = for ((j, z) ← r3)
        where (y.tasks = j)
          return z))))
```

We can now state our key correctness property: evaluating shredded queries and stitching the results back together yields the same results as evaluating the original nested query directly.

THEOREM 4. *If $\vdash L : \text{Bag } A$ then:*

$$\text{stitch}(\mathcal{H}[L]_{\text{Bag } A}) = \mathcal{N}[L]$$

PROOF SKETCH. We omit the full proof due to space limits; it is available in the full version of this paper. The proof introduces several intermediate definitions. Specifically, we consider an *annotated semantics* for nested queries in which each collection element is tagged with an index, and we show that this semantics is consistent with the ordinary semantics if the annotations are erased. We then prove the correctness of shredding and stitching with respect to the annotated semantics, and the desired result follows. \square

6. INDEXING SCHEMES

So far, we have worked with canonical indexes of the form $a \diamond 1.2.3$. These could be represented in SQL by using multiple columns (padding with NULLs if necessary) since for a given query the length of the dynamic part is bounded by the number of for-comprehensions in the query. This imposes space and running time overhead due to constructing and maintaining the indexes. Instead, in this section we consider alternative, more compact indexing schemes.

We can define alternative indexing schemes by varying the *index* parameter of the shredded semantics (see Section 5.1). Not all possible instantiations are valid. To identify those that are, we first define a function for computing the canonical indexes of a nested query result.

$$\begin{aligned} \mathcal{I}[L] &= \mathcal{I}[L]_{\varepsilon, 1} \\ \mathcal{I}[\langle \ell_i = C_i \rangle_{i=1}^n]_{\rho, \iota} &= \text{concat}([\mathcal{I}[C_i]_{\rho, \iota}]_{i=1}^n) \\ \mathcal{I}[\langle \ell_i = M_i \rangle_{i=1}^n]_{\rho, \iota} &= \text{concat}([\mathcal{I}[M_i]_{\rho, \iota}]_{i=1}^n) \\ \mathcal{I}[X]_{\rho, \iota} &= [] \\ \mathcal{I}[\text{for } ([x_i \leftarrow \ell_i]_{i=1}^n \text{ where } X) \text{ return }^a M]_{\rho, \iota} &= \\ \text{concat}([a \diamond \iota.j :: \mathcal{I}[M]_{\rho[x_i \mapsto r_i]_{i=1}^n, \iota.j}} \\ | \langle j, \vec{r} \rangle \leftarrow \text{enum}(\{\vec{r} \mid [r_i \leftarrow \mathcal{I}[\ell_i]_{\rho[x_i \mapsto r_i]_{i=1}^n, \iota.j}}])]) \end{aligned}$$

$$\begin{aligned}
\mathcal{S}[\llbracket L \rrbracket] &= \mathcal{S}[\llbracket L \rrbracket]_{\varepsilon,1} & \mathcal{S}[\langle \ell = N \rangle_{i=1}^n]_{\rho,\iota} &= \langle \ell_i = \mathcal{S}[\llbracket N_i \rrbracket]_{\rho,\iota} \rangle_{i=1}^n & \mathcal{S}[\llbracket a \diamond \text{out} \rrbracket]_{\rho,\iota,i} &= \text{index}(a \diamond \iota) \\
& & \mathcal{S}[\llbracket X \rrbracket]_{\rho,\iota} &= \mathcal{N}[\llbracket X \rrbracket]_{\rho} & \mathcal{S}[\llbracket a \diamond \text{in} \rrbracket]_{\rho,\iota,i} &= \text{index}(a \diamond \iota.i) \\
\mathcal{S}[\llbracket \biguplus_{i=1}^n C_i \rrbracket]_{\rho,\iota} &= \text{concat}([\mathcal{S}[\llbracket C_i \rrbracket]_{\rho,\iota}]_{i=1}^n) & & & \mathcal{S}[\text{return}^a N]_{\rho,\iota} &= [\mathcal{S}[\llbracket N \rrbracket]_{\rho,\iota}] \\
\mathcal{S}[\text{for } ([x_i \leftarrow t_i]_{i=1}^n \text{ where } X) C]_{\rho,\iota} &= \text{concat}([\mathcal{S}[\llbracket C \rrbracket]_{\rho[x_i \mapsto r_i]_{i=1}^n, \iota, j} \mid \langle j, \vec{r} \rangle \leftarrow \text{enum}(\vec{r} \mid [r_i \leftarrow \llbracket t_i \rrbracket]_{i=1}^n, \mathcal{N}[\llbracket X \rrbracket]_{\rho[x_i \mapsto r_i]_{i=1}^n})])
\end{aligned}$$

Figure 5: Semantics of shredded queries

An indexing function $\text{index} : \text{Index} \rightarrow X$ is *valid* with respect to the closed nested query L if it is injective and defined on every canonical index in $\mathcal{I}[\llbracket L \rrbracket]$. The only requirement on indexes in the proof of Theorem 4 is that index is valid. We consider two alternative valid indexing schemes: *natural* and *flat* indexes.

6.1 Natural indexes

Natural indexes are synthesised from row data. In order to generate a natural index for a query every table must have a key, that is, a collection of fields guaranteed to be unique for every row in the table. For sets, this is always possible by using all of the field values as a key; this idea is used in Van den Bussche’s simulation for sets [31]. However, for bags this is not always possible, so using natural indexes may require adding extra key fields.

Given a table t , let key_t be the function that given a row r of t returns the key fields of r . We now define a function to compute the list of natural indexes for a query L .

$$\begin{aligned}
\mathcal{I}^\natural[\llbracket L \rrbracket] &= \mathcal{I}^\natural[\llbracket L \rrbracket]_{\varepsilon} \\
\mathcal{I}^\natural[\llbracket \biguplus_{i=1}^n C_i \rrbracket]_{\rho} &= \text{concat}([\mathcal{I}^\natural[\llbracket C_i \rrbracket]_{\rho}]_{i=1}^n) \\
\mathcal{I}^\natural[\langle \ell_i = M_i \rangle_{i=1}^n]_{\rho} &= \text{concat}([\mathcal{I}^\natural[\llbracket M_i \rrbracket]_{\rho}]_{i=1}^n) \\
\mathcal{I}^\natural[\llbracket X \rrbracket]_{\rho} &= [] \\
\mathcal{I}^\natural[\text{for } ([x_i \leftarrow t_i]_{i=1}^n \text{ where } X) \text{return}^a M]_{\rho} &= \\
&\text{concat}([a \diamond \langle \text{key}_{t_i}(r_i) \rangle_{i=1}^n :: \mathcal{I}^\natural[\llbracket M \rrbracket]_{\rho[x_i \mapsto r_i]_{i=1}^n} \\
&\mid [r_i \leftarrow \llbracket t_i \rrbracket]_{i=1}^n, \mathcal{N}[\llbracket X \rrbracket]_{\rho[x_i \mapsto r_i]_{i=1}^n}])
\end{aligned}$$

If $a \diamond \iota$ is the i -th element of $\mathcal{I}[\llbracket L \rrbracket]$, then $\text{index}_L^\natural(a \diamond \iota)$ is defined as the i -th element of $\mathcal{I}^\natural[\llbracket L \rrbracket]$. The natural indexing scheme is defined by setting $\text{index} = \text{index}_L^\natural$.

An advantage of natural indexes is that they can be implemented in plain SQL, so for a given comprehension all where clauses can be amalgamated (using the \wedge operator) and no auxiliary subqueries are needed. The downside is that the type of a dynamic index may still vary across the component comprehensions of a shredded query, complicating implementation of the query (due to the need to pad some subqueries with null columns) and potentially decreasing performance due to increased data movement. We now consider an alternative, in which `row_number` is used to generate dynamic indexes.

6.2 Flat indexes and let-insertion

The idea of flat indexes is to enumerate all of the canonical dynamic indexes associated with each static index and use the enumeration as the dynamic index.

Let ι be the i -th element of the list $\{\iota' \mid a \diamond \iota' \leftarrow \mathcal{I}[\llbracket L \rrbracket]\}$, then $\text{index}_L^\flat(a \diamond \iota) = \langle a, i \rangle$. The flat indexing scheme is defined by setting $\text{index} = \text{index}_L^\flat$. Let $\mathcal{I}^\flat[\llbracket L \rrbracket] = [\text{index}_L^\flat(I) \mid I \leftarrow \mathcal{I}[\llbracket L \rrbracket]]$ and let $\mathcal{S}^\flat[\llbracket L \rrbracket]$ be $\mathcal{S}[\llbracket L \rrbracket]$ where $\text{index} = \text{index}_L^\flat$.

In this section, we give a translation called let-insertion that uses let-binding and an index primitive to manage flat indexes. In the next section, we take the final step from this language to SQL.

Our semantics for shredded queries uses canonical indexes. We now specify a target language providing flat indexes. In order to do so, we introduce let-bound sub-queries, and translate each compre-

hension into the following form:

$$\begin{aligned}
&\text{let } q = \text{for } (\vec{G}_{\text{out}} \text{ where } X_{\text{out}}) \text{return } N_{\text{out}} \text{ in} \\
&\text{for } (\vec{G}_{\text{in}} \text{ where } X_{\text{in}}) \text{return } N_{\text{in}}
\end{aligned}$$

The special index expression is available in each loop body, and is bound to the current index value.

Following let-insertion, the types are as before, except indexes are represented as pairs of integers $\langle a, d \rangle$ where a is the static component and d is the dynamic component.

$$\begin{aligned}
\text{Types} \quad A, B &::= \text{Bag } \langle \langle \text{Int}, \text{Int} \rangle, F \rangle \\
\text{Flat types} \quad F &::= O \mid \langle \vec{\ell} : \vec{F} \rangle \mid \langle \text{Int}, \text{Int} \rangle
\end{aligned}$$

The syntax of terms is adapted as follows:

$$\begin{aligned}
\text{Query terms} \quad L, M &::= \biguplus \vec{C} \\
\text{Comprehensions} \quad C &::= \text{let } q = S \text{ in } S' \\
\text{Subqueries} \quad S &::= \text{for } (\vec{G} \text{ where } X) \text{return } N \\
\text{Data sources} \quad u &::= t \mid q \\
\text{Generators} \quad G &::= x \leftarrow u \\
\text{Inner terms} \quad N &::= X \mid R \mid \text{index} \\
\text{Record terms} \quad R &::= \langle \vec{\ell} = \vec{N} \rangle \\
\text{Base terms} \quad X &::= x.\vec{\ell} \mid c(\vec{X}) \mid \text{empty } L
\end{aligned}$$

The semantics of let-inserted queries is given in Figure 6. Rather than maintaining a canonical index, it generates a flat index for each subquery.

We first give the translation on shredded types as follows:

$$\begin{aligned}
\mathbf{L}(O) &= O \\
\mathbf{L}(\langle \vec{\ell} : \vec{F} \rangle) &= \langle \vec{\ell} : \mathbf{L}(\vec{F}) \rangle \\
\mathbf{L}(\text{Index}) &= \langle \text{Int}, \text{Int} \rangle \\
\mathbf{L}(\text{Bag } \langle \text{Index}, F \rangle) &= \text{Bag } \langle \langle \text{Int}, \text{Int} \rangle, \mathbf{L}(F) \rangle
\end{aligned}$$

For example:

$$\mathbf{L}(A_2) = \text{Bag } \langle \langle \text{Int}, \text{Int} \rangle, \langle \text{name} : \text{String}, \text{tasks} : \langle \text{Int}, \text{Int} \rangle \rangle \rangle$$

Without loss of generality we rename all the bound variables in our source query to ensure that all bound variables have distinct names, and that none coincides with the distinguished name z used for let-bindings. The let-insertion translation \mathbf{L} is defined in Figure 7, where we use the following auxiliary functions:

$$\begin{aligned}
\text{expand}(x, t) &= \langle \ell_i = x.\ell_i \rangle_{i=1}^n \\
&\quad \text{where } \Sigma(t) = \text{Bag } \langle \vec{\ell} : \vec{A} \rangle \\
\text{gens } (\text{for } (\vec{G} \text{ where } X) C) &= \vec{G} :: \text{gens } C \\
\text{gens } (\text{return}^a N) &= [] \\
\text{conds } (\text{for } (\vec{G} \text{ where } X) C) &= X :: \text{conds } C \\
\text{conds } (\text{return}^a N) &= [] \\
\text{body } (\text{for } (\vec{G} \text{ where } X) C) &= \text{body } C \\
\text{body } (\text{return}^a N) &= N
\end{aligned}$$

Each comprehension is rearranged into two sub-queries. The first generates the outer indexes. The second computes the results. The translation sometimes produces n -ary projections in order to refer to values bound by the first subquery inside the second.

$$\begin{aligned}
\mathcal{L}[\llbracket L \rrbracket] &= \mathcal{L}[\llbracket L \rrbracket]_\varepsilon & \mathcal{L}[\llbracket t \rrbracket]_\rho &= \llbracket t \rrbracket & \mathcal{L}[\langle \ell_j = N_j \rangle_{j=1}^m]_{\rho, i} &= \langle \ell_j = \mathcal{L}[\llbracket N_j \rrbracket]_{\rho, i} \rangle_{j=1}^m & \mathcal{L}[\llbracket \text{index} \rrbracket]_{\rho, i} &= i \\
\mathcal{L}[\llbracket \biguplus_{j=1}^m C_j \rrbracket]_\rho &= \text{concat}([\mathcal{L}[\llbracket C_j \rrbracket]_\rho]_{j=1}^m) & \mathcal{L}[\llbracket q \rrbracket]_\rho &= \rho(q) & \mathcal{L}[\llbracket X \rrbracket]_{\rho, i} &= \mathcal{N}[\llbracket X \rrbracket]_\rho \\
\mathcal{L}[\llbracket \text{let } q = S_{\text{out}} \text{ in } S_{\text{in}} \rrbracket]_\rho &= \mathcal{L}[\llbracket S_{\text{in}} \rrbracket]_{\rho[q \mapsto \mathcal{L}[\llbracket S_{\text{out}} \rrbracket]_\rho]} & & & & & & \\
\mathcal{L}[\llbracket \text{for } ([x_j \leftarrow u_j]_{j=1}^m \text{ where } X) \text{ return } N \rrbracket]_\rho &= [\mathcal{L}[\llbracket N \rrbracket]_{\rho[x_j \mapsto r_j]_{j=1}^m, i} \mid \langle i, \vec{r} \rangle \leftarrow \text{enum}([\vec{r} \mid [r_j \leftarrow \mathcal{L}[\llbracket u_j \rrbracket]_{\rho, j=1}^m, \mathcal{N}[\llbracket X \rrbracket]_{\rho[x_j \mapsto r_j]_{j=1}^m}])}]
\end{aligned}$$

Figure 6: Semantics of let-inserted shredded queries

$$\begin{aligned}
\mathbf{L}(\biguplus_{i=1}^n C_i) &= \biguplus_{i=1}^n \mathbf{L}(C_i) \\
\mathbf{L}(C) &= \text{let } q = (\text{for } (\vec{G}_{\text{out}} \text{ where } X_{\text{out}}) \text{ return } \langle R_{\text{out}}, \text{index} \rangle) \text{ in for } (z \leftarrow q, \vec{G}_{\text{in}} \text{ where } \mathbf{L}_{\vec{y}}(X_{\text{in}})) \text{ return } \mathbf{L}_{\vec{y}}(N) \\
\text{where } \vec{G}_{\text{out}} &= \text{concat}(\text{init}(\text{gens } C)) & y = t = \vec{G}_{\text{out}} & \vec{G}_{\text{in}} = \text{last}(\text{gens } C) & N = \text{body } C \\
X_{\text{out}} &= \bigwedge \text{init}(\text{conds } C) & R_{\text{out}} = \langle \text{expand}(y_i, t_i) \rangle_{i=1}^n & X_{\text{in}} = \text{last}(\text{conds } C) & n = \text{length } \vec{G}_{\text{out}} \\
\mathbf{L}_{\vec{y}}(x.\ell) &= \begin{cases} x.\ell, & \text{if } x \notin \{y_1, \dots, y_n\} \\ z.1.i.\ell, & \text{if } x = y_i \end{cases} & \mathbf{L}_{\vec{y}}(\text{empty } L) &= \text{empty } (\mathbf{L}_{\vec{y}}(L)) \\
\mathbf{L}_{\vec{y}}(c(X_1, \dots, X_m)) &= c(\mathbf{L}_{\vec{y}}(X_1), \dots, \mathbf{L}_{\vec{y}}(X_m)) & \mathbf{L}_{\vec{y}}(\biguplus_{i=1}^n C_i) &= \biguplus_{i=1}^n \mathbf{L}_{\vec{y}}(C_i) \\
\mathbf{L}_{\vec{y}}(\langle \ell_j = X_j \rangle_{j=1}^m) &= \langle \ell_j = \mathbf{L}_{\vec{y}}(X_j) \rangle_{j=1}^m & \mathbf{L}_{\vec{y}}(\text{for } (\vec{G} \text{ where } X) \text{ return }^a \langle a, N \rangle) &= \text{for } (\vec{G} \text{ where } \mathbf{L}_{\vec{y}}(X)) \\
& & & \text{return } \langle a, \mathbf{L}_{\vec{y}}(N) \rangle \\
\mathbf{L}_{\vec{y}}(a \diamond d) &= \langle a, \mathbf{L}(d) \rangle & \mathbf{L}(\text{out}) &= z.2 \\
& & \mathbf{L}(\text{in}) &= \text{index}
\end{aligned}$$

Figure 7: The let-insertion translation

For example, applying \mathbf{L} to q_1 from Section 4.2 yields:

```

for (x ← departments)
  return ⟨⟨T, 1⟩, ⟨dept = x.name, people = index⟩⟩

```

and q_2 becomes:

```

(let q = for (x ← departments) return ⟨⟨dept = x.name⟩, index⟩ in
  for (z ← q, y ← employees) where (z.1.1.name = y.dept ∧
    (y.salary < 1000 ∨ y.salary > 1000000))
  return ⟨⟨⟨a, z.2⟩, ⟨name = y.name, tasks = ⟨b, index⟩⟩⟩⟩)
⊔
(let q = for (x ← departments) return ⟨⟨dept = x.name⟩, index⟩ in
  for (z ← q, y ← contacts) where (z.1.1.name = y.dept ∧ y.client)
  return ⟨⟨⟨a, z.2⟩, ⟨name = y.name, tasks = ⟨d, index⟩⟩⟩⟩)

```

As a sanity check, we show that the translation is type-preserving:

THEOREM 5. *Given shredded query $\vdash M : \text{Bag } \langle \text{Index}, F \rangle$, then $\vdash \mathbf{L}(M) : \mathbf{L}(\text{Bag } \langle \text{Index}, F \rangle)$.*

To prove the correctness of let-insertion, we need to show that the shredded semantics and let-inserted semantics agree. In the statement of the correctness theorem, recall that $\mathcal{S}^b[\llbracket L \rrbracket]$ refers to the version of $\mathcal{S}[\llbracket L \rrbracket]$ using $\text{index} = \text{index}_L^b$.

THEOREM 6. *Suppose $\vdash L : A$ and $\llbracket L \rrbracket_\rho = M$. Then $\mathcal{S}^b[\llbracket M \rrbracket] = \mathcal{L}[\llbracket \mathbf{L}(M) \rrbracket]$.*

PROOF SKETCH. The high-level idea is to separate results into data and indexes and compare each separately. It is straightforward, albeit tedious, to show that the different definitions are equal if we replace all dynamic indexes by a dummy value. It then remains to show that the dynamic indexes agree. The pertinent case is the translation of a comprehension:

$$[\text{for } (\vec{G}_i \leftarrow X_i)]_{i=1}^n \text{ for } (\vec{G}_{\text{in}} \leftarrow X_{\text{in}}) \text{ return }^b \langle a \diamond \text{out}, N \rangle$$

which becomes $\text{let } q = S_{\text{out}} \text{ in } S_{\text{in}}$ for suitable S_{out} and S_{in} . The dynamic indexes computed by S_{out} coincide exactly with those of $\mathcal{T}^b[\llbracket L \rrbracket]$ at static index a , and the dynamic indexes computed by S_{in} , if there are any, coincide exactly with those of $\mathcal{T}^b[\llbracket L \rrbracket]$. \square

7. CONVERSION TO SQL

Earlier translation stages have employed nested records for convenience, but SQL does not support nested records. At this stage,

we eliminate nested records from queries. For example, we can represent a nested record $\langle a = \langle b = 1, c = 2 \rangle, d = 3 \rangle$ as a flat record $\langle a_b = 1, a_c = 2, d = 3 \rangle$. The (standard) details are presented in Appendix E.

In order to interpret shredded, flattened, let-inserted terms as SQL, we interpret index generators using SQL's OLAP facilities.

Query terms	$L ::= (\text{union all}) \vec{C}$
Comprehensions	$C ::= \text{with } q \text{ as } (S) C \mid S'$
Subqueries	$S ::= \text{select } R \text{ from } \vec{G} \text{ where } X$
Data sources	$u ::= t \mid q$
Generators	$G ::= u \text{ as } x$
Inner terms	$N ::= X \mid \text{row_number}() \text{ over } (\text{order by } \vec{X})$
Record terms	$R ::= \vec{N} \text{ as } \vec{\ell}$
Base terms	$X ::= x.\ell \mid c(\vec{X}) \mid \text{empty } L$

Continuing our example, $\mathbf{L}(q_1)$ and $\mathbf{L}(q_2)$ translate to q'_1 and q'_2 where:

```

q'_1 = select x.name as i1_name,
  row_number() over (order by x.name) as i1_people
  from departments as x
q'_2 = (with q as (select x.name as i1_name,
  row_number() over (order by x.name) as i2
  from departments as x)
  select a as i1_1, z.i2 as i1_2, y.name as i2_name, b as i2_tasks_1,
  row_number() over (order by z.i1_name, z.i2,
    y.dept, y.employee, y.salary) as i2_tasks_2
  from employees as y, q as z
  where (z.i1_name = y.dept ∧
    (y.salary < 1000 ∨ y.salary > 1000000)))
union all
(with q as (select x.name as i1_name,
  row_number() over (order by x.name) as i2
  from departments as x)
  select a as i1_1, z.i2 as i1_2, y.name as i2_name, d as i2_tasks_1,
  row_number() over (order by z.i1_name, z.i2,
    y.dept, y.name, y.client) as i2_tasks_2
  from contacts as y, q as z
  where (z.i1_name = y.dept ∧ y.client))

```

Modulo record flattening, the above fragment of SQL is almost isomorphic to the image of the let-insertion translation. The only significant difference is the use of `row_number` in place of `index`. Each instance of `index` in the body R of a subquery of the form

```

QF1: SELECT e.emp FROM employees e
      WHERE e.salary > 10000
QF2: SELECT e.emp, t.task
      FROM employees e, tasks t
      WHERE e.emp = t.emp
QF3: SELECT e1.emp, e2.emp
      FROM employees e1, employees e2
      WHERE e1.dpt = e2.dpt
          AND e1.salary = e2.salary
          AND e1.emp <> e2.emp
QF4: (SELECT t.emp FROM tasks t
      WHERE t.task = 'abstract')
      UNION ALL (SELECT e.emp FROM employees
                  WHERE e.salary > 50000)
QF5: (SELECT t.emp FROM tasks t
      WHERE t.task = 'abstract')
      MINUS
      (SELECT e.emp FROM employees e
       WHERE e.salary > 50000)
QF6: ((SELECT t.emp FROM tasks t
       WHERE t.task = 'abstract')
      UNION ALL (SELECT e.emp FROM employees e
                  WHERE e.salary > 50000))
      MINUS
      ((SELECT t.emp FROM tasks t
       WHERE t.task = 'enthuse')
      UNION ALL (SELECT e.emp FROM employees e
                  WHERE e.salary > 10000))

```

Figure 8: SQL versions of flat queries used in experiments

for $(\overrightarrow{x} \leftarrow \overrightarrow{t} \text{ where } X)$ return R is simulated by a term of the form $\text{row_number}() \text{ over } (\text{order by } \overrightarrow{x}.\overrightarrow{\ell})$, where:

$$\begin{aligned}
x_i &: \langle \ell_{i,1} : A_{i,1}, \dots, \ell_{i,m_i} : A_{i,m_i} \rangle \\
\overrightarrow{x}.\overrightarrow{\ell} &= x_1.\ell_{1,1}, \dots, x_1.\ell_{1,m_1}, \dots, x_n.\ell_{n,1}, \dots, x_n.\ell_{n,m_n}
\end{aligned}$$

A possible concern is that row_number is non-deterministic. It computes row numbers ordered by the supplied columns, but if there is a tie, then it is free to order the equivalent rows in any order. However, we always order by *all* columns of all tables referenced from the current subquery, so our use of row_number is always deterministic. (An alternative could be to use nonstandard features such as PostgreSQL’s `OID` or MySQL’s `ROWNUM`, but sorting would still be necessary to ensure consistency across inner and outer queries.)

8. EXPERIMENTAL EVALUATION

Ferry’s loop-lifting translation has been implemented in Links previously by Ulrich, a member of the Ferry team [30], following the approach described by Grust et al. [12] to generate SQL:1999 algebra plans, then calling Pathfinder [14] to optimise and evaluating the resulting SQL on a PostgreSQL database. We have also implemented query shredding in Links, running against PostgreSQL; our implementation¹ does not use Pathfinder. We performed initial experiments with a larger ad hoc query benchmark, and developed some optimisations, including inlining certain `WITH` clauses to unblock rewrites, using keys for row numbering, and implementing stitching in one pass to avoid construction of intermediate in-memory data structures that are only used once and then discarded. We report on shredding with all of these optimisations enabled.

Benchmark queries. There is no standard benchmark for queries returning nested results. In particular, the popular TPC-

¹<http://github.com/slindley/links/tree/shredding>

```

Q1 : for (d ← departments)
      return ((name = d.name,
               employees = employeesOfDept d,
               contacts = contactsOfDept d))
Q2 : for (d ← Q1)
      where (all d.employees (λx.contains x.task "abstract"))
      return (dept = d.dept)
Q3 : for (e ← employees)
      return (name = e.name, task = tasksOfEmp(e))
Q4 : for (d ← departments)
      return (dept = d.dept, employees = for (e ← employees)
                                             where (d.dept = e.dept)
                                             return e.employee)
Q5 : for (t ← tasks) return (a = t.task, b = employeesByTask t)
Q6 : for (d ← Q1)
      return ((department = d.name,
               people =
                 getTasks(outliers(d.employees)) (λy. y.tasks)
                 ⊔ getTasks(clients(d.contacts)) (λy. return "buy"))))

```

Figure 9: Nested queries used in experiments

H benchmark is not suitable for comparing shredding and loop-lifting: the TPC-H queries do not return nested results, and can be run directly on any SQL-compliant RDBMS, so neither approach needs to do any work to produce an SQL query.

We selected twelve queries over the organisation schema described in Section 3 to use as a benchmark. The first six queries, named QF1–QF6, return flat results and can be translated to SQL using existing techniques, without requiring either shredding or loop-lifting. We considered these queries as a sanity check and in order to measure the overhead introduced by loop-lifting and shredding. Figure 8 shows the SQL versions of these queries.

We also selected six queries Q1–Q6 that do involve nesting, either within query results or in intermediate stages. They are shown in Figure 9; they use the auxiliary functions defined in Section 3. Q1 is the query Q_{org} that builds the nested organisation view from Section 3. Q2 is a flat query that computes a flat result from Q1 consisting of all departments in which all employees can do the “abstract” task; it is a typical example of a query that employs higher-order functions. Q3 returns records containing each employee and the list of tasks they can perform. Q4 returns records containing departments and the set of employees in each department. Q5 returns a record of tasks paired up with sets of employees and their departments. Q6 is the outliers query Q introduced in Section 3.

Experimental results. We measured the query execution time for all queries on randomly generated data, where we vary the number of departments in the organisation from 4 to 4096 (by powers of 2). Each department has on average 100 employees and each employee has 0–2 tasks, and the largest (4096 department) database was approximately 500MB. Although the test datasets are moderate in size, they suffice to illustrate the asymptotic trends in the comparative performance of the different techniques. All tests were performed using PostgreSQL 9.2 running on a MacBook Pro with 4-core 2.6GHz CPU, 8GB RAM and 500GB SSD storage, with the database server running on the same machine (hence, negligible network latency). We measure *total* time to translate a nested query to SQL, evaluate the resulting SQL queries, and stitch the results together to form a nested value, measured from Links.

We evaluated each query using query shredding and loop-lifting, and for the flat queries we also measured the time for Links’ default

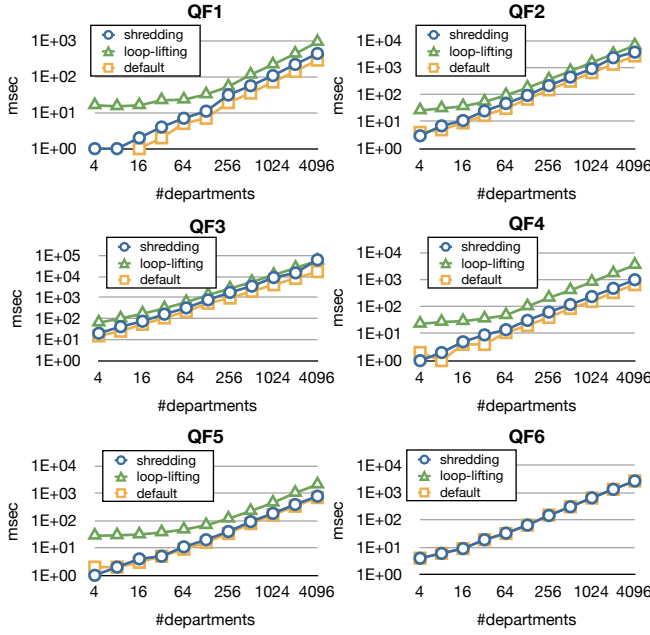


Figure 10: Experimental results (flat queries)

(flat) query evaluation. The experimental results for the flat queries are shown in Figure 10 and for the nested queries in Figure 11. Note that both axes are logarithmic. All times are in milliseconds; the times are medians of 5 runs. The times for small data sizes provide a comparison of the overhead associated with shredding, loop-lifting or Links’ default query normalisation algorithm.

Discussion. We should re-emphasise that Ferry (and Ulrich’s loop-lifting implementation for Links) supports grouping and aggregation features that are not handled by our translation. We focused on queries that both systems can handle, but Ferry has a clear advantage for grouping and aggregation queries or when ordering is important (e.g. sorting or top- k queries). Ferry is based on list semantics, while our approach handles multiset semantics. So, some performance differences may be due to our exploitation of multiset-based optimisations that Ferry (by design) does not exploit.

The results for flat queries show that shredding has low per-query overhead in most cases compared to Links’ default flat query evaluation, but the queries it generates are slightly slower. Loop-lifting, on the other hand, has a noticeable per-query overhead, likely due to its invocation of Pathfinder and associated serialisation costs. In some cases, such as QF4 and QF5, loop-lifting is noticeably slower asymptotically; this appears to be due to additional sorting needed to maintain list semantics. We encountered a bug that prevented loop-lifting from running on query QF6; however, shredding had negligible overhead for this query. In any case, the overhead of either shredding or loop-lifting for flat queries is irrelevant: we can simply evaluate such queries using Links’ default flat query evaluator. Nevertheless, these results show that the overhead of shredding for such queries is not large, suggesting that the queries it generates are similar to those currently generated by Links. (Manual inspection of the generated queries confirms this.)

The results for nested queries are more mixed. In most cases, the overhead of loop-lifting is dominant on small data sizes, which again suggests that shredding may be preferable for OLTP or Web workloads involving rapid, small queries. Loop-lifting scales poorly

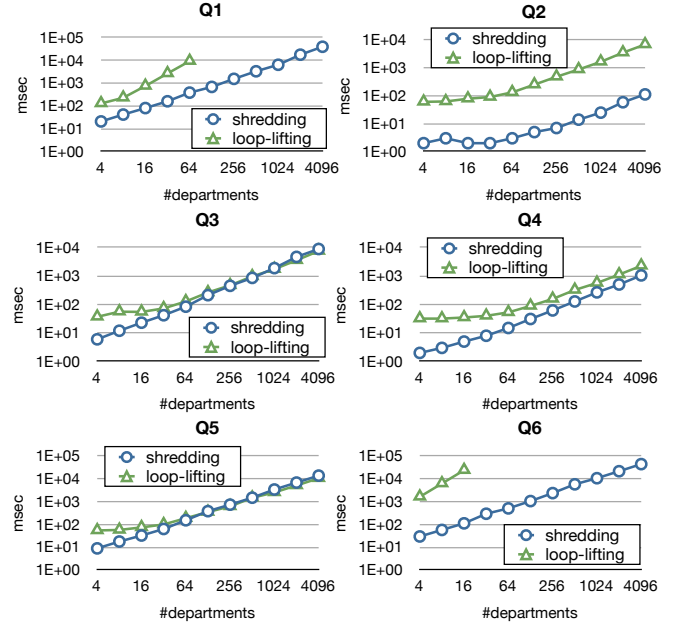


Figure 11: Experimental results (nested queries)

on two queries (Q1 and Q6), and did not finish within 1 minute even for small data sizes. Both Q1 and Q6 involve 3 levels of nesting and in the innermost query, loop-lifting generates queries with Cartesian products inside OLAP operators such as DENSE_RANK or ROW_NUMBER that Pathfinder was not able to remove. The queries generated by shredding in these cases avoid this pathological behaviour. For other queries, such as Q2 and Q4, loop-lifting performs better but is still slower than shredding as data size increases. Comparison of the queries generated by loop-lifting and shredding reveals that loop-lifting encountered similar problems with hard-to-optimize OLAP operations. Finally, for Q3 and Q5, shredding is initially faster (due to the overhead of loop-lifting and calling Pathfinder) but as data size increases, loop-lifting wins out. Inspection of these generated queries reveals that the queries themselves are similar, but the shredded queries involve more data movement. Also, loop-lifting returns sorted results, so it avoids in-memory hashing or sorting while constructing the nested result. It should be possible to incorporate similar optimisations into shredding to obtain comparable performance.

Our experiments show that shredding performs similarly or better than loop-lifting on our (synthetic) benchmark queries on moderate (up to 500MB) databases. Further work may need to be done to investigate scalability to larger databases or consider more realistic query benchmarks.

9. RELATED AND FUTURE WORK

We have discussed related work on nested queries, Links, Ferry and LINQ in the introduction. Besides Cooper [7], several authors have recently considered higher-order query languages. Benedikt et al. [2, 32] study higher-order queries over flat relations. The Links approach was adapted to LINQ in F# by Cheney et al. [6]. Higher-order features are also being added to XQuery 3.0 [25].

Research on shredding XML data into relations and evaluating XML queries over such representations [18] is superficially similar to our work in using various indexing or numbering schemes to handle nested data. Grust et al.’s work on translating XQuery

to SQL via Pathfinder [14] is a mature solution to this problem, and Grust et al. [13] discuss optimisations in the presence of unordered data processing in XQuery. However, XQuery's ordered data model and treatment of node identity would block transformations in our algorithm that assume unordered, pure operations.

We can now give a more detailed comparison of our approach with the indexing strategies in Van den Bussche's work and in Ferry. Van den Bussche's approach uses natural indexes (that is, n -tuples of ids), but does not preserve multiset semantics. Our approach preserves multiplicity and can use natural indexes, we also propose a flat indexing scheme based on `row_number`. In Ferry's indexing scheme, the surrogate indexes only link adjacent nesting levels, whereas our indexes take information at all higher levels into account. Our flat indexing scheme relies on this property, and Ferry's approach does not seem to be an instance of ours (or vice versa). Ferry can generate multiple SQL:1999 operations and Pathfinder tries to merge them but cannot always do so. Our approach generates `row_number` operations only at the end, and does not rely on Pathfinder. Finally, our approach uses normalisation and tags parts of the query to disambiguate branches of unions.

Loop-lifting has been implemented in Links by Ulrich [30], and Grust and Ulrich [16] recently presented techniques for supporting higher-order functions as query results. By using Ferry's loop-lifting translation and Pathfinder, Ulrich's system also supports list semantics and aggregation and grouping operations; to our knowledge, it is an open problem to either prove their correctness or adapt these techniques to fit our approach. Ferry's approach supports a list-based semantics for queries, while we assume a bag-based semantics (matching SQL's default behaviour). Either approach can accommodate set-based semantics simply by eliminating duplicates in the final result. In fact, however, we believe the core query shredding translation (Sections 4–6) works just as well for a list semantics. The only parts that rely on unordered semantics are normalisation (Section 2.2) and conversion to SQL (Section 7). We leave these extensions to future work.

Our work is also partly inspired by work on unnesting for nested data parallelism. Blelloch and Sabot [3] give a compilation scheme for NESL, a data-parallel language with nested lists; Suciu and Tannen [27] give an alternative scheme for a nested list calculus. This work may provide an alternative (and parallelisable) implementation strategy for Ferry's list-based semantics [12].

10. CONCLUSION

Combining efficient database access with high-level programming abstractions is challenging in part because of the limitations of flat database queries. Currently, programmers must write flat queries and manually convert the results to a nested form. This damages declarativity and maintainability. Query shredding can help to bridge this gap. Although it is known from prior work that query shredding is possible in principle, and some implementations (such as Ferry) support this, getting the details right is tricky, and can lead to queries that are not optimised well by the relational engine. Our contribution is an alternative shredding translation that handles queries over bags and delays use of OLAP operations until the final stage. Our translation compares favourably to loop-lifting in performance, and should be straightforward to extend and to incorporate into other language-integrated query systems.

Acknowledgements. We thank Ezra Cooper, Torsten Grust, and Jan Van den Bussche for comments. We are very grateful to Alex Ulrich for sharing his implementation of Ferry in Links and for extensive discussions. This work was supported by EPSRC grants EP/J014591/1 and EP/K034413/1 and a Google Research Award.

11. REFERENCES

- [1] A. Beckmann. Exact bounds for lengths of reductions in typed λ -calculus. *Journal of Symbolic Logic*, 66:1277–1285, 2001.
- [2] M. Benedikt, G. Puppis, and H. Vu. Positive higher-order queries. In *PODS*, pages 27–38, 2010.
- [3] G. E. Blelloch and G. W. Sabot. Compiling collection-oriented languages onto massively parallel computers. *J. Parallel Distrib. Comput.*, 8:119–134, February 1990.
- [4] P. Buneman, L. Libkin, D. Suciu, V. Tannen, and L. Wong. Comprehension syntax. *SIGMOD Record*, 23, 1994.
- [5] P. Buneman, S. A. Naqvi, V. Tannen, and L. Wong. Principles of programming with complex objects and collection types. *Theor. Comput. Sci.*, 149(1):3–48, 1995.
- [6] J. Cheney, S. Lindley, and P. Wadler. A practical theory of language-integrated query. In *ICFP*, pages 403–416. ACM, 2013.
- [7] E. Cooper. The script-writer's dream: How to write great SQL in your own language, and be sure it will succeed. In *DBPL*, pages 36–51, 2009.
- [8] E. Cooper, S. Lindley, P. Wadler, and J. Yallop. Links: web programming without tiers. In *FMCO*, volume 4709 of *LNCS*, 2007.
- [9] P. de Groote. On the strong normalisation of intuitionistic natural deduction with permutation-conversions. *Inf. Comput.*, 178(2):441–464, 2002.
- [10] L. Fegaras and D. Maier. Optimizing object queries using an effective calculus. *ACM Trans. Database Syst.*, 25:457–516, December 2000.
- [11] T. Grust, M. Mayr, J. Rittinger, and T. Schreiber. Ferry: Database-supported program execution. In *SIGMOD*, June 2009.
- [12] T. Grust, J. Rittinger, and T. Schreiber. Avalanche-safe LINQ compilation. *PVLDB*, 3(1), 2010.
- [13] T. Grust, J. Rittinger, and J. Teubner. eXrQuery: Order indifference in XQuery. In *ICDE*, pages 226–235, 2007.
- [14] T. Grust, J. Rittinger, and J. Teubner. Pathfinder: XQuery off the relational shelf. *IEEE Data Eng. Bull.*, 31(4):7–14, 2008.
- [15] T. Grust, S. Sakr, and J. Teubner. XQuery on SQL hosts. In *VLDB*, pages 252–263, 2004.
- [16] T. Grust and A. Ulrich. First-class functions for first-order database engines. In *DBPL*, 2013.
- [17] W. Kim. On optimizing an SQL-like nested query. *ACM Trans. Database Syst.*, 7(3):443–469, Sept. 1982.
- [18] R. Krishnamurthy, R. Kaushik, and J. F. Naughton. XML-SQL query translation literature: The state of the art and open problems. In *Xsym*, pages 1–18, 2003.
- [19] S. Lindley. Extensional rewriting with sums. In *TLCA*, 2007.
- [20] S. Lindley and J. Cheney. Row-based effect types for database integration. In *TLDI*, 2012.
- [21] S. Lindley and I. Stark. Reducibility and $\top\top$ -lifting for computation types. In *TLCA*, 2005.
- [22] E. Meijer, B. Beckman, and G. M. Bierman. LINQ: reconciling object, relations and XML in the .NET framework. In *SIGMOD*, 2006.
- [23] C. Olston, B. Reed, U. Srivastava, R. Kumar, and A. Tomkins. Pig latin: a not-so-foreign language for data processing. In *SIGMOD*, pages 1099–1110, 2008.
- [24] B. Pierce. *Types and Programming Languages*. MIT Press, 2002.
- [25] J. Robie, D. Chamberlin, M. Dyck, and J. Snelson. XQuery 3.0: An XML query language. W3C Proposed Recommendation, October 2013. <http://www.w3.org/TR/xquery-30/>.
- [26] H. J. Schek and M. H. Scholl. The relational model with relation-valued attributes. *Inf. Syst.*, 11:137–147, April 1986.
- [27] D. Suciu and V. Tannen. Efficient compilation of high-level data parallel algorithms. In *SPAA*, pages 57–66. ACM, 1994.
- [28] D. Syme, A. Granicz, and A. Cisternino. *Expert F# 3.0*. Apress, 2012.
- [29] W. W. Tait. Intensional interpretations of functionals of finite type I. *Journal of Symbolic Logic*, 32(2):198–212, June 1967.
- [30] A. Ulrich. A Ferry-based query backend for the Links programming language. Master's thesis, University of Tübingen, 2011. Code supplement: <http://github.com/slindley/links/tree/sand>.
- [31] J. Van den Bussche. Simulation of the nested relational algebra by

the flat relational algebra, with an application to the complexity of evaluating powerset algebra expressions. *Theor. Comput. Sci.*, 254(1-2), 2001.

- [32] H. Vu and M. Benedikt. Complexity of higher-order queries. In *ICDT*, pages 208–219. ACM, 2011.
- [33] P. Wadler. Comprehending monads. *Math. Struct. in Comp. Sci.*, 2(4), 1992.
- [34] L. Wong. Normal forms and conservative extension properties for query languages over collection types. *J. Comput. Syst. Sci.*, 52(3), 1996.
- [35] L. Wong. Kleisli, a functional query system. *J. Funct. Programming*, 10(1), 2000.

APPENDIX

A. BEHAVIOUR OF VAN DEN BUSSCHE’S SIMULATION ON MULTISETS

As noted in the Introduction, Van den Bussche’s simulation of nested queries via flat queries does not work properly over multisets. To illustrate the problem, consider a simple query $R \cup S$, where R and S have the same schema $\text{Bag} \langle A : \text{Int}, B : \text{Bag Int} \rangle$. Suppose R and S have the following values:

	<table><tr><td>A</td><td>B</td></tr><tr><td>1</td><td>{1}</td></tr><tr><td>2</td><td>{2}</td></tr></table>	A	B	1	{1}	2	{2}		<table><tr><td>A</td><td>B</td></tr><tr><td>1</td><td>{3,4}</td></tr><tr><td>2</td><td>{2}</td></tr></table>	A	B	1	{3,4}	2	{2}
A	B														
1	{1}														
2	{2}														
A	B														
1	{3,4}														
2	{2}														

then their multiset union is:

$R \cup S =$	<table><tr><th>A</th><th>B</th></tr><tr><td>1</td><td>{1}</td></tr><tr><td>1</td><td>{3, 4}</td></tr><tr><td>2</td><td>{2}</td></tr><tr><td>2</td><td>{2}</td></tr></table>	A	B	1	{1}	1	{3, 4}	2	{2}	2	{2}
A	B										
1	{1}										
1	{3, 4}										
2	{2}										
2	{2}										

Van den Bussche’s simulation (like ours) represents these nested values by queries over flat tables, such as the following:

$R_1 =$	<table><tr><th>A</th><th>id</th></tr><tr><td>1</td><td>a</td></tr><tr><td>2</td><td>b</td></tr></table>	A	id	1	a	2	b	$R_2 =$	<table><tr><th>id</th><th>B</th></tr><tr><td>a</td><td>1</td></tr><tr><td>b</td><td>2</td></tr></table>	id	B	a	1	b	2		
A	id																
1	a																
2	b																
id	B																
a	1																
b	2																
$S_1 =$	<table><tr><th>A</th><th>id</th></tr><tr><td>1</td><td>a</td></tr><tr><td>2</td><td>b</td></tr></table>	A	id	1	a	2	b	$S_2 =$	<table><tr><th>id</th><th>B</th></tr><tr><td>a</td><td>3</td></tr><tr><td>a</td><td>4</td></tr><tr><td>b</td><td>2</td></tr></table>	id	B	a	3	a	4	b	2
A	id																
1	a																
2	b																
id	B																
a	3																
a	4																
b	2																

where a, b are arbitrary distinct ids. Note however that R and S have overlapping ids, so if we simply take the union of R_1 and S_1 , and of R_2 and S_2 respectively, we will get:

$Wrong_1 =$	<table><tr><th>A</th><th>id</th></tr><tr><td>1</td><td>a</td></tr><tr><td>2</td><td>b</td></tr><tr><td>1</td><td>a</td></tr><tr><td>2</td><td>b</td></tr></table>	A	id	1	a	2	b	1	a	2	b	$Wrong_2 =$	<table><tr><th>id</th><th>B</th></tr><tr><td>a</td><td>1</td></tr><tr><td>b</td><td>2</td></tr><tr><td>a</td><td>3</td></tr><tr><td>a</td><td>4</td></tr><tr><td>b</td><td>2</td></tr></table>	id	B	a	1	b	2	a	3	a	4	b	2
	A	id																							
	1	a																							
	2	b																							
	1	a																							
2	b																								
id	B																								
a	1																								
b	2																								
a	3																								
a	4																								
b	2																								

corresponding to nested value:

$Wrong =$	<table><tr><th>A</th><th>B</th></tr><tr><td>1</td><td>{1, 3, 4}</td></tr><tr><td>1</td><td>{1, 3, 4}</td></tr><tr><td>2</td><td>{2, 2}</td></tr><tr><td>2</td><td>{2, 2}</td></tr></table>	A	B	1	{1, 3, 4}	1	{1, 3, 4}	2	{2, 2}	2	{2, 2}
A	B										
1	{1, 3, 4}										
1	{1, 3, 4}										
2	{2, 2}										
2	{2, 2}										

Instead, both Van den Bussche’s simulation and our approach avoid clashes among ids when taking unions. Van den Bussche’s simulation does this by adding two new id fields to represent the result of a union, say id_1 and id_2 . The tuples originating from R will have *equal* id_1 and id_2 values, while those originating from S will have *different* id_1 and id_2 values. In order to do this, the simulation defines two queries, one for each result table. The first query is of the form:

$$\begin{aligned} T_1 &= (R_1 \times (id_1 : x, id_2 : x) \mid x \in \mathbf{adom}) \\ &\cup (S_1 \times (id_1 : x, id_2 : x') \mid x \neq x' \in \mathbf{adom}) \end{aligned}$$

and similarly

$$\begin{aligned} T_2 &= (R_2 \times (id_1 : x, id_2 : x) \mid x \in \mathbf{adom}) \\ &\cup (S_2 \times (id_1 : x, id_2 : x') \mid x \neq x' \in \mathbf{adom}) \end{aligned}$$

where **adom** is the active domain of the database (in this case, **adom** = {1, 2, 3, 4, a, b}) — this is of course also definable as a query. Thus, in this example, the result is of the form

$T_1 =$	<table border="1"> <thead> <tr> <th>A</th><th>id</th><th>id_1</th><th>id_2</th></tr> </thead> <tbody> <tr><td>1</td><td>a</td><td>x</td><td>x</td></tr> <tr><td>2</td><td>b</td><td>y</td><td>y</td></tr> <tr><td>1</td><td>a</td><td>z</td><td>z'</td></tr> <tr><td>2</td><td>b</td><td>v</td><td>v'</td></tr> </tbody> </table>	A	id	id_1	id_2	1	a	x	x	2	b	y	y	1	a	z	z'	2	b	v	v'	$T_1 =$	<table border="1"> <thead> <tr> <th>id</th><th>id_1</th><th>id_2</th><th>B</th></tr> </thead> <tbody> <tr><td>a</td><td>x</td><td>x</td><td>1</td></tr> <tr><td>b</td><td>y</td><td>y</td><td>2</td></tr> <tr><td>a</td><td>z</td><td>z'</td><td>3</td></tr> <tr><td>a</td><td>w</td><td>w'</td><td>4</td></tr> <tr><td>b</td><td>v</td><td>v'</td><td>2</td></tr> </tbody> </table>	id	id_1	id_2	B	a	x	x	1	b	y	y	2	a	z	z'	3	a	w	w'	4	b	v	v'	2
A	id	id_1	id_2																																												
1	a	x	x																																												
2	b	y	y																																												
1	a	z	z'																																												
2	b	v	v'																																												
id	id_1	id_2	B																																												
a	x	x	1																																												
b	y	y	2																																												
a	z	z'	3																																												
a	w	w'	4																																												
b	v	v'	2																																												

where x, y are any elements of **adom** (rows mentioning x and y stand for 6 instances) and $z \neq z', w \neq w'$ and $v \neq v'$ are any pairs of distinct elements of **adom** (rows mentioning these variables stand for 30 instances of distinct pairs from **adom**). This leads to an $O(|\mathbf{adom}| * |R| + |\mathbf{adom}|^2 * |S|)$ blowup in the number of tuples. Specifically, for our example, $|T_1| = 72$, whereas the actual number of tuples in a natural representation of $R \cup S$ is only 9. In a set semantics, the set value simulated by these tables is correct even with all of the extra tuples; however, for a multiset semantics, this quadratic blowup is not correct — even for our example, with $|\mathbf{adom}| = 6$, the result of evaluating $R \cup S$ yields a different number of tuples from the result of evaluating $S \cup R$, and neither represents the correct multiset in an obvious way. It may be possible (given knowledge of the query and active domain, but not the source database) to “decode” the flat query results and obtain the correct nested result, but doing so appears no easier than developing an alternative, direct algorithm.

B. TYPING RULES

The (standard) typing rules for λ_{NRC} queries are shown in Figure 12. The typing rules for shredded terms, packages, and values are shown in Figures 13, 14, 16 and 18.

C. QUERY NORMALISATION

Following previous work [20], we separate query normalisation into a rewriting phase and a type-directed structurally recursive function. Where we diverge from this previous work is that we extend the rewrite relation to hoist all conditionals up to the nearest comprehension (in order to simplify the rest of the development), and the structurally recursive function is generalised to handle nested data. Thus normalisation can be divided into three stages.

- The first stage performs symbolic evaluation, that is, β -reduction and commuting conversions, flattening unnecessary nesting and eliminating higher-order functions.
- The second stage hoists all conditionals up to the nearest enclosing comprehension in order to allow them to be converted to where clauses.
- The third and final stage hoists all unions up to the top-level, η -expands tables and variables, and turns all conditionals into where clauses.

Weak normalisation and strong normalisation. Given a term M and a rewrite relation \rightsquigarrow , we write $M \not\rightsquigarrow$ if M is irreducible, that is no rewrite rules in \rightsquigarrow apply to M . The term M is said to be in *normal form*.

A term M is *weakly normalising* with respect to a rewrite relation \rightsquigarrow_r , or *r-WN*, if there exists a finite reduction sequence

$$M \rightsquigarrow M_1 \rightsquigarrow \dots \rightsquigarrow M_n \not\rightsquigarrow$$

A term M is *strongly normalising* with respect to a rewrite relation \rightsquigarrow_r , or *r-SN*, if every reduction sequence starting from M is finite.

If M is *r-SN*, then we write $\max_r(M)$ for the maximum length of a reduction sequence starting from M .

A rewrite relation \rightsquigarrow_r is *weakly-normalising*, or *WN*, if all terms M are weakly-normalising with respect to \rightsquigarrow_r . Similarly, a rewrite relation \rightsquigarrow_r is *strongly-normalising*, or *SN*, if all terms M are strongly-normalising with respect to \rightsquigarrow_r .

For any \rightsquigarrow_r we can define a partial function nf_r such that $nf_r(M)$ is a normal form of M , if one exists, otherwise $nf_r(M)$ is undefined. Note that there may be more than one normal form for a given M ; nf_r chooses a particular normal form for each M that has one. If \rightsquigarrow_r is weakly normalising, then nf_r is a total function.

We now describe each normalisation stage in turn.

C.1 Symbolic evaluation

The β -rules perform symbolic evaluation, including substituting argument values for function parameters, record field projection, conditionals where the test is known to be true or false, or iteration over singleton bags.

$$\begin{aligned}
&(\lambda x. N) M \rightsquigarrow_c N[x := M] \\
&\langle \ell = M \rangle. \ell_i \rightsquigarrow_c M_i \\
&\text{if true then } M \text{ else } N \rightsquigarrow_c M \\
&\text{if false then } M \text{ else } N \rightsquigarrow_c N \\
&\text{for } (x \leftarrow \text{return } M) N \rightsquigarrow_c N[x := M]
\end{aligned}$$

Fundamentally, β -rules always follow the same pattern. Each is associated with a particular type constructor T , and the left-hand side always consists of an *introduction* form for T inside an *elimination* form for T . For instance, in the case of functions (the first β -rule above), the introduction form is a lambda and the elimination form is an application. Applying a β -rule *eliminates* T in the sense that the introduction form from the left-hand side either no longer appears or has been replaced by a term of a simpler type on the right-hand side.

Each instance of a β -rule is associated with an *elimination frame*. An *elimination frame* is simply the elimination form with a designated *hole* $[\]$ that can be *plugged* with another expression. For instance, elimination frames for the function β -rule are of the form $E[\] = [\] M$. If we plug an introduction form $\lambda x. N$ into $E[\]$, written $E[\lambda x. N]$, then we obtain the left-hand side of the associated β -rule $(\lambda x. N) M$.

(This notion of an expression with a hole that can be filled by another expression is commonly used in rewriting and operational semantics for programming languages [24, ch. 19]; here, it is not essential but helps cut down the number of explicit rules, and helps highlight commonality between rules.) The elimination frames of λ_{NRC} are as follows.

$$E[\] ::= [\] M \mid [\].\ell \mid \text{if } [\] \text{ then } M \text{ else } N \mid \text{for } (x \leftarrow [\]) N$$

The following rules express that comprehensions, conditionals, empty bag constructors, and unions can always be *hoisted* out of the above elimination frames. In the literature such rules are often called *commuting conversions*. They are necessary in order to expose all possible β -reductions. For instance,

$$(\text{if } M \text{ then } \langle \ell = N \rangle \text{ else } N').\ell$$

cannot β -reduce, but if M then $\langle \ell = N \rangle.\ell$ else $N'.\ell$ can.

$$\begin{aligned}
&E[\text{for } (x \leftarrow M) N] \rightsquigarrow_c \text{for } (x \leftarrow M) E[N] \\
&E[\text{if } L \text{ then } M \text{ else } N] \rightsquigarrow_c \text{if } L \text{ then } E[M] \text{ else } E[N] \\
&E[\emptyset] \rightsquigarrow_c \emptyset \\
&E[M_1 \uplus M_2] \rightsquigarrow_c E[M_1] \uplus E[M_2]
\end{aligned}$$

For example:

$$(\text{if } L \text{ then } M_1 \text{ else } M_2) M \rightsquigarrow_c \text{if } L \text{ then } M_1 M \text{ else } M_2 M$$

$\frac{\text{VAR}}{\Gamma, x : A \vdash x : A}$	$\frac{\text{CONSTANT} \quad \Sigma(c) = \langle O_1, \dots, O_n \rangle \rightarrow O' \quad [\Gamma \vdash M_i : O_i]_{i=1}^n}{\Gamma \vdash c(M_1, \dots, M_n) : O'}$	$\frac{\text{LAM} \quad \Gamma, x : A \vdash M : B}{\Gamma \vdash \lambda x. M : A \rightarrow B}$
$\frac{\text{APP} \quad \Gamma \vdash M : A \rightarrow B \quad \Gamma \vdash N : A}{\Gamma \vdash M N : B}$	$\frac{\text{RECORD} \quad [\Gamma \vdash M_i : A_i]_{i=1}^n}{\Gamma \vdash \langle \ell_i = M_i \rangle_{i=1}^n : \langle \ell_i = A_i \rangle_{i=1}^n}$	$\frac{\text{PROJECT} \quad \Gamma \vdash M : \langle \ell_i : A_i \rangle_{i=1}^n}{\Gamma \vdash M.\ell_j : A_j}$
$\frac{\text{IF} \quad \Gamma \vdash M : \text{Bool} \quad \Gamma \vdash N : A \quad \Gamma \vdash N' : A}{\Gamma \vdash \text{if } M \text{ then } N \text{ else } N' : A}$	$\frac{\text{EMPTY}}{\Gamma \vdash \emptyset : \text{Bag } A}$	$\frac{\text{SINGLETON} \quad \Gamma \vdash M : A}{\Gamma \vdash \text{return } M : \text{Bag } A}$
$\frac{\text{FOR} \quad \Gamma \vdash M : \text{Bag } A \quad \Gamma, x : A \vdash N : \text{Bag } B}{\Gamma \vdash \text{for } (x \leftarrow M) N : \text{Bag } B}$	$\frac{\text{ISEMPTY} \quad \Gamma \vdash M : \text{Bag } A}{\Gamma \vdash \text{empty } M : \text{Bool}}$	$\frac{\text{TABLE} \quad \Sigma(t) = \text{Bag } \langle \overrightarrow{\ell : \vec{O}} \rangle}{\Gamma \vdash \text{table } t : \text{Bag } \langle \overrightarrow{\ell : \vec{O}} \rangle}$
$\frac{\text{UNION} \quad \Gamma \vdash M : \text{Bag } A \quad \Gamma \vdash N : \text{Bag } A}{\Gamma \vdash M \uplus N : \text{Bag } A}$		

Figure 12: Typing rules for higher-order nested queries

$\frac{\text{VAR}}{\Gamma, x : \langle \overrightarrow{\ell : \vec{F}} \rangle \vdash x : \langle \overrightarrow{\ell : \vec{F}} \rangle}$	$\frac{\text{CONSTANT} \quad \Sigma(c) = \langle O_1, \dots, O_n \rangle \rightarrow O' \quad [\Gamma \vdash X_i : O_i]_{i=1}^n}{\Gamma \vdash c(X_1, \dots, X_n) : O'}$	$\frac{\text{RECORD} \quad [\Gamma \vdash N_i : A_i]_{i=1}^n}{\Gamma \vdash \langle \ell_i = N_i \rangle_{i=1}^n : \langle \ell_i = A_i \rangle_{i=1}^n}$
$\frac{\text{PROJECT} \quad \Gamma \vdash x : \langle \ell_i : A_i \rangle_{i=1}^n}{\Gamma \vdash x.\ell_j : A_j}$	$\frac{\text{INDEX}}{\Gamma \vdash a \diamond d : \text{Index}}$	$\frac{\text{SINGLETON} \quad \Gamma \vdash I : \text{Index} \quad \Gamma \vdash N : F}{\Gamma \vdash \text{return}^a \langle I, N \rangle : \text{Bag } \langle \text{Index}, F \rangle}$
$\frac{\text{FOR} \quad [\Sigma(t_i) = \text{Bag } A_i]_{i=1}^n \quad \Gamma, [x_i : A_i]_{i=1}^n \vdash X : \text{Bool} \quad \Gamma, [x_i : A_i]_{i=1}^n \vdash C : \text{Bag } \langle \text{Index}, F \rangle}{\Gamma \vdash \text{for } ([x_i \leftarrow t_i]_{i=1}^n \text{ where } X) C : \text{Bag } \langle \text{Index}, F \rangle}$		$\frac{\text{UNION} \quad [\Gamma \vdash C_i : \text{Bag } \langle \text{Index}, F \rangle]_{i=1}^n}{\Gamma \vdash \biguplus \vec{C} : \text{Bag } \langle \text{Index}, F \rangle}$
		$\frac{\text{ISEMPTY} \quad \Gamma \vdash L : \text{Bag } \langle \text{Index}, F \rangle}{\Gamma \vdash \text{empty } L : \text{Bool}}$

Figure 13: Typing rules for shredded terms

$\frac{\text{BASE}}{\vdash O(\Lambda_S) : O(\mathcal{T}_S)}$	$\frac{\text{RECORD} \quad \vdash \hat{A}_i(\Lambda_S) : \hat{A}'_i(\mathcal{T}_S)]_{i=1}^n}{\vdash \langle \ell_i : \hat{A}_i(\Lambda_S) \rangle_{i=1}^n : \langle \ell_i : \hat{A}'_i(\mathcal{T}_S) \rangle_{i=1}^n}$	$\frac{\text{BAG} \quad \vdash \hat{A}(\Lambda_S) : \hat{A}'(\mathcal{T}_S) \quad \vdash L : A \quad \text{erase}(\hat{A}(\Lambda_S)) = A}{\vdash (\text{Bag } \hat{A}(\Lambda_S))^L : (\text{Bag } \hat{A}'(\mathcal{T}_S))^A}$
--	--	--

We write Λ_S for the set of shredded terms, and \mathcal{T}_S for the set of shredded types.

Figure 14: Typing rules for shredded packages

Note that some combinations of elimination frames and rewrite rule are impossible in a well-typed term, such as if \emptyset then M else N . For the purposes of reduction we treat empty like an uninterpreted constant, that is, we do reduce inside emptiness tests, but they do not in any other way interact with the reduction rules.

Next we prove that \sim_c is strongly normalising. The proof is based on a previous proof of strong normalisation for simply-typed λ -calculus with sums [19], which generalises the $\top\top$ -lifting approach [21], which in turn extends Tait's proof of strong normalisation for simply-typed λ -calculus [29].

Frame stacks.

$$\begin{aligned} \text{(frame stacks)} \quad & S ::= Id \mid S \circ E \\ \text{(stack length)} \quad & |Id| = 0 \\ & |S \circ E| = |S| + 1 \\ \text{(plugging)} \quad & Id[M] = M \\ & (S \circ E)[M] = S[(E[M])] \end{aligned}$$

Following previous work [19] we assume variables are annotated with types. We assume variables are annotated with types. We write $A \multimap B$ for the type of frame stack S , if $S[M] : B$ for all terms $M : A$.

Frame stack reduction.

$$S \sim_c S' \stackrel{\text{def}}{\iff} \forall M. S[M] \sim_c S'[M] \\ \iff S[x] \sim_c S'[x]$$

Frame stacks are closed under reduction. A frame stack S is *c-strongly normalising*, or *c-SN*, if all reduction sequences starting from S are finite.

LEMMA 7. *If $S \sim_c S'$, for frame stacks S, S' , then $|S'| \leq |S|$.*

PROOF. Induction on the structure of S . \square

Reducibility. We define *reducibility* as follows:

- Id is reducible.
- $S \circ ([] N) : (A \rightarrow B) \multimap C$ is reducible if S and N are reducible.
- $S \circ ([] \ell) : (\ell : \vec{A}) \multimap C$ is reducible if S is reducible.
- $S : \text{Bag } A \multimap C$ is reducible if $S[\text{return } M]$ is *c-SN* for all reducible $M : A$.
- $S : \text{Bool} \multimap C$ is reducible if $S[\text{true}]$ is *c-SN* and $S[\text{false}]$ is *c-SN*.
- $M : A$ is reducible if $S[M]$ is *c-SN* for all reducible $S : A \multimap C$.

LEMMA 8. *If $M : A$ is reducible then M is *c-SN*.*

PROOF. Follows immediately from reducibility of Id and the definition of reducibility on terms. \square

LEMMA 9. *$x : A$ is reducible.*

PROOF. By induction on A using Lemma 7 and Lemma 8. \square

COROLLARY 10. *If $S : A \multimap C$ is reducible then S is *c-SN*.*

Each type constructor has an associated β -rule. Each β -rule gives rise to an SN-closure property.

LEMMA 11 (SN-CLOSURE).

(\rightarrow) *If $S[M[x : = N]]$ and N are *c-SN* then $S[(\lambda x.M) N]$ is *c-SN*.*

$(\langle \rangle)$ *If \vec{M} , are *c-SN* then $S[\langle \ell = \vec{M} \rangle . \ell_i]$ is *c-SN*.*

$(\text{Bag } \rightarrow)$ *If $S[N[x : = M]]$ and M are *c-SN* then $S[\text{for } (x \leftarrow \text{return } M) N]$ is *c-SN*.*

(Bool) *If $S[N]$ and $S[N']$ are *c-SN* then $S[\text{if true then } N \text{ else } N']$ is *c-SN* and $S[\text{if false then } N \text{ else } N']$ is *c-SN*.*

PROOF.

(\rightarrow) : By induction on $\max_c(S) + \max_c(M) + \max_c(N)$.

$(\langle \rangle)$: By induction on

$$\max_c S + \left(\sum_{i=1}^n \max_c(M_i) \right) + \max_c(N) + \left(\sum_{i=1}^n \max_c(M'_i) \right).$$

$(\text{Bag } \rightarrow)$: By induction on

$$|S| + \max_c(S[N[x : = M]]) + \max_c(M).$$

(Bool) : By induction on $|S| + \max_c(S[N]) + \max_c(S[N'])$.

This completes the proof. \square

Now we obtain reducibility-closure properties for each type constructor.

LEMMA 12 (REDUCIBILITY-CLOSURE).

(\rightarrow) *If $M[x : = N]$ is reducible for all reducible N , then $\lambda x.M$ is reducible.*

$(\langle \rangle)$ *If \vec{M} are reducible, then $\langle \ell = \vec{M} \rangle$ is reducible.*

$(\text{Bag } \rightarrow)$ *If M is reducible, $N[x : = M']$ is reducible for all reducible M' , then for $(x \leftarrow \text{return } M) N$ is reducible.*

(Bool) *If M, N, N' are reducible then if M then N else, N' is reducible.*

PROOF. Each property follows from the corresponding part of Lemma 11 using Lemma 8 and Corollary 10. \square

We also require additional closure properties for the empty bag and union constructs.

LEMMA 13 (REDUCIBILITY-CLOSURE II).

(\emptyset) *The empty bag \emptyset is reducible.*

(\uplus) *If M, N are reducible, then $M \uplus N$ is reducible.*

PROOF.

(\emptyset) : Suppose $S : \text{Bag } A \multimap C$ is reducible. We need to prove that $S[\emptyset]$ is *c-SN*. The proof is by induction on $|S| + \max_c(S)$. The only interesting case is hoisting the empty bag out of a bag elimination frame, which simply decreases the size of the frame stack by 1.

(\uplus) : Suppose $M, N : \text{Bag } A$, and $S : \text{Bag } A \multimap C$ are reducible. We need to show that $S[M \uplus N]$ is *c-SN*. The proof is by induction on $|S| + \max_c(S[M]) + \max_c(S[N])$. The only interesting case is hoisting the union out of a bag elimination frame, which again decreases the size of the frame stack by 1, whilst leaving the other components of the induction measure unchanged.

This completes the proof. \square

THEOREM 14. *Let M be any term. Suppose $x_1 : A_1, \dots, x_n : A_n$ includes all the free variables of M . If $N_1 : A_1, \dots, N_n : A_n$ are reducible then $M[\vec{x} = \vec{N}]$ is reducible.*

PROOF. By induction on the structure of terms using Lemma 12 and Lemma 13. \square

THEOREM 15 (STRONG NORMALISATION). *The relation \sim_c is strongly normalising.*

PROOF. Let M be a term with free variables \vec{x} . By Lemma 9, \vec{x} are reducible. Hence, by Theorem 14, M is c -SN. \square

It is well known that β -reduction in simply-typed λ -calculus has non-elementary complexity in the worst case [1]. The relation \sim_c includes β -reduction, so it must be at least as bad (we conjecture that it has the same asymptotic complexity, as \sim_c can be reduced to β -reduction on simply-typed λ -calculus via a CPS translation following de Groote [9]). However, we believe that the asymptotic complexity is unlikely to pose a problem in practice, as the kind of higher-order code that exhibits worst-case behaviour is rare. It has been our experience with Links that query normalisation time is almost always dominated by SQL execution time.

C.2 If hoisting

To hoist conditionals (if-expressions) out of constant applications, records, unions, and singleton bag constructors, we define if-hoisting frames as follows:

$$F[] ::= c(\vec{M}, [], \vec{N}) \mid \langle \ell' = \vec{M}, \ell = [], \ell'' = \vec{N} \rangle \\ \mid [] \uplus N \mid M \uplus [] \mid \text{return } []$$

The if-hoisting rule says that if an expression contains an if-hoisting frame around a conditional, then we can lift the conditional up and push the frame into both branches:

$$F[\text{if } L \text{ then } M \text{ else } N] \sim_h \text{if } L \text{ then } F[M] \text{ else } F[N]$$

We write $\text{size}(M)$ for the size of M , as in the total number of syntax constructors in M .

LEMMA 16.

1. If M, N, N' are h -SN then $\text{if } M \text{ then } N \text{ else } N'$ is h -SN.
2. If M, N are h -SN then $M \uplus N$ is h -SN.
3. If \vec{M} are h -SN then $c(\vec{M})$ is h -SN.
4. If \vec{M} are h -SN then $\langle \ell = \vec{M} \rangle$ is h -SN.

PROOF.

1: By induction on $\langle \text{max}_h(M), \text{size}(M), \text{max}_h(N) + \text{max}_h(N'), \text{size}(N) + \text{size}(N') \rangle$.

2: By induction on $\langle \text{max}_h(M) + \text{max}_h(N), \text{size}_h(M) + \text{size}_h(N) \rangle$ using (1).

3 and 4: By induction on

$$\langle \sum_{i=1}^n \text{max}_h(M_i), \sum_{i=1}^n \text{size}(M_i) \rangle$$

using (1).

This concludes the proof. \square

PROPOSITION 17. *The relation \sim_h is strongly normalising.*

PROOF. By induction on the structure of terms using Lemma 16. \square

C.3 The query normalisation function

The following definition of the function norm generalises the normalisation algorithm from previous work [20].

$$\text{norm}_A(M) = \langle \text{nf}_h(\text{nf}_c(M)) \rangle_A$$

where:

$$\begin{aligned} \langle c([X_i : O_i]_{i=1}^n) \rangle_O &= c(\langle [X_i]_{i=1}^n \rangle_{O_i}) \\ \langle x.\ell \rangle_O &= x.\ell \\ \langle \text{empty}(M : \text{Bag } A) \rangle_{\text{Bool}} &= \text{empty}(\langle M \rangle_{\text{Bag } A}) \\ \langle M \rangle_{\langle \ell_i : A_i \rangle_{i=1}^n} &= \langle \ell_i = F \langle M \rangle_{A_i, \ell_i} \rangle_{i=1}^n \\ \langle M \rangle_{\text{Bag } A} &= \biguplus (\mathcal{B} \langle M \rangle_{A, [], \text{true}}) \\ \mathcal{B} \langle \text{return } M \rangle_{A, \vec{G}, L}^* &= [\text{for } (\vec{G} \text{ where } L) \text{ return } \langle M \rangle_A] \\ \mathcal{B} \langle \text{for } (x \leftarrow t) M \rangle_{A, \vec{G}, L}^* &= \mathcal{B} \langle M \rangle_{A, \vec{G} + [x \leftarrow t], L}^* \\ \mathcal{B} \langle \text{table } t \rangle_{A, \vec{G}, L}^* &= \mathcal{B} \langle \text{return } x \rangle_{A, \vec{G} + [x \leftarrow t], L}^* \\ &\quad (x \text{ fresh}) \\ \mathcal{B} \langle \emptyset \rangle_{A, \vec{G}, L}^* &= [] \\ \mathcal{B} \langle M \uplus N \rangle_{A, \vec{G}, L}^* &= \mathcal{B} \langle M \rangle_{A, \vec{G}, L}^* \uplus \mathcal{B} \langle N \rangle_{A, \vec{G}, L}^* \\ \mathcal{B} \langle \text{if } L' M N \rangle_{A, \vec{G}, L}^* &= \mathcal{B} \langle M \rangle_{A, \vec{G}, L \wedge L'}^* \\ &\quad \uplus \mathcal{B} \langle N \rangle_{A, \vec{G}, L \wedge \neg L'}^* \\ \mathcal{F} \langle x \rangle_{A, \ell_i} &= \langle x.\ell_i \rangle_A \\ \mathcal{F} \langle \ell_i = M_i \rangle_{i=1}^n_{A, \ell_i} &= \langle M_i \rangle_A \end{aligned}$$

Strictly speaking, in order for the above definition of norm_A to make sense we need the two rewrite relations to be confluent. It is easily verified that the relation \sim_c is locally confluent, and hence by strong normalisation and Newman's Lemma it is confluent. The relation \sim_h is not confluent as the ordering of hoisting determines the final order in which booleans are eliminated. However, it is easily seen to be confluent modulo reordering of conditionals, which in turn means that norm_A is well-defined if we identify terms modulo commutativity of conjunction, which is perfectly reasonable given that conjunction is indeed commutative.

THEOREM 18. *The function norm_A terminates.*

PROOF. The result follows immediately from strong normalisation of \sim_c and \sim_h , and the fact that the functions $\langle - \rangle$, $\mathcal{B} \langle - \rangle^*$, and $\mathcal{F} \langle - \rangle$ are structurally recursive (modulo expanding out the right-hand-side of the definition of $\mathcal{B} \langle \text{table } t \rangle_{A, \vec{G}, L}^*$). \square

D. PROOF OF CORRECTNESS OF SHRED-DING

The main correctness result is that if we shred an annotated normalised nested query, run all the resulting shredded queries, and stitch the shredded results back together, then that is the same as running the nested query directly.

Intuitively, the idea is as follows:

1. define shredding on nested values;
2. show that shredding commutes with query execution; and
3. show that stitching is a left-inverse to shredding of values.

Unfortunately, this strategy is a little too naive to work, because nested values do not contain enough information about the structure of the source query in order to generate the same indexes that are generated by shredded queries.

To fix the problem, we will give an annotated semantics $\mathcal{A} \llbracket - \rrbracket$ that is defined only on queries that have first been converted to normal form and annotated with static indexes as described in Section 4. We will ensure that $\mathcal{N} \llbracket \text{erase}(L) \rrbracket = \text{erase}(\mathcal{A} \llbracket L \rrbracket)$, where we overload the erase function to erase annotations on terms and values. The structure of the resulting proof is depicted in the diagram in Figure 15, where we view a query as a function from the in-

RESULT	RECORD	CONSTANT
$\frac{[\vdash v_i : A_i]_{i=1}^n \quad [\vdash J : Index]_{i=1}^n}{\vdash [v_i @ J_i]_{i=1}^n : \text{Bag } A}$	$\frac{[\vdash v_i : A_i]_{i=1}^n}{\vdash \langle \ell_i = v_i \rangle_{i=1}^n : \langle \ell_i = A_i \rangle_{i=1}^n}$	$\frac{\Sigma(c) = O}{\vdash c : O}$

Figure 16: Typing rules for nested annotated values

$$\begin{aligned}
\mathcal{S}[\![L]\!] &= \mathcal{S}[\![L]\!]_{\varepsilon,1} & \mathcal{S}[\![\ell = N]_{i=1}^n]\!]_{\rho,\iota} &= \langle \ell_i = \mathcal{S}[\![N_i]\!]_{\rho,\iota} \rangle_{i=1}^n & \mathcal{S}[\![a \diamond \text{out}]\!]_{\rho,\iota,i} &= \text{index}(a \diamond \iota) \\
& & \mathcal{S}[\![X]\!]_{\rho,\iota} &= \mathcal{N}[\![X]\!]_{\rho} & \mathcal{S}[\![a \diamond \text{in}]\!]_{\rho,\iota,i} &= \text{index}(a \diamond \iota.i) \\
\mathcal{S}[\![\biguplus_{i=1}^n C_i]\!]_{\rho,\iota} &= \text{concat}([\mathcal{S}[\![C_i]\!]_{\rho,\iota}]_{i=1}^n) & \mathcal{S}[\![\text{return}^a N]\!]_{\rho,\iota} &= [\mathcal{S}[\![N]\!]_{\rho,\iota} @ \text{index}(a \diamond \iota)] \\
\mathcal{S}[\![\text{for } ([x_i \leftarrow t_i]_{i=1}^n \text{ where } X) C]\!]_{\rho,\iota} &= \text{concat}([\mathcal{S}[\![C]\!]_{\rho[x_i \mapsto r_i]_{i=1}^n, \iota, j} \mid \langle j, \vec{r} \rangle \leftarrow \text{enum}([\vec{r} \mid [r_i \leftarrow \llbracket t_i \rrbracket]_{i=1}^n, \mathcal{N}[\![X]\!]_{\rho[x_i \mapsto r_i]_{i=1}^n})])])
\end{aligned}$$

Figure 17: Semantics of shredded queries, with annotations




Figure 15: Correctness of shredding and stitching

terpretation of the input schema Σ to the interpretation of its result type. To prove correctness is to prove that this diagram commutes.

D.1 Annotated semantics of nested queries

We annotate bag elements with distinct indexes as follows:

Results	$s ::= [w_1 @ I_1, \dots, w_m @ I_m]$
Inner values	$w ::= c \mid r \mid s$
Rows	$r ::= \langle \ell_1 = w_1, \dots, \ell_n = w_n \rangle$
Indexes	I, J

Typing rules for these values are in Figure 16. The index annotations $@I$ on collection elements are needed solely for our correctness proof, and do not need to be present at run time.

The *canonical* choice for representing indexes is to take $I = a \diamond \iota$. We also allow for alternative indexing schemes in by parameterising the semantics over an indexing function *index* mapping each canonical index $a \diamond \iota.j$ to a concrete representation. In the simplest case, we take *index* to be the identity function (where static indexes are viewed as distinct integers). We will later consider other definitions of *index* that depend on the particular query being shredded. Informally, the only constraints on *index* are that it is defined on every canonical index $a \diamond \iota$ that we might need in order to shred a query, and it is injective on all canonical indexes. We will formalise these constraints later, but for now, *index* can be assumed to be the identity function.

Given an indexing function *index*, the semantics of nested queries

on annotated values is defined as follows:

$$\begin{aligned}
\mathcal{A}[\![L]\!] &= \mathcal{A}[\![L]\!]_{\varepsilon,1} \\
\mathcal{A}[\![\biguplus_{i=1}^n C_i]\!]_{\rho,\iota} &= \text{concat}([\mathcal{A}[\![C_i]\!]_{\rho,\iota}]_{i=1}^n) \\
\mathcal{A}[\![\ell_i = M_i]_{i=1}^n]\!]_{\rho,\iota} &= \langle \ell_i = \mathcal{A}[\![M_i]\!]_{\rho,\iota} \rangle_{i=1}^n \\
\mathcal{A}[\![X]\!]_{\rho,\iota} &= \mathcal{N}[\![X]\!]_{\rho} \\
\mathcal{A}[\![\text{for } ([x_i \leftarrow t_i]_{i=1}^n \text{ where } X) \text{return}^a M]\!]_{\rho,\iota} &= \\
&[\mathcal{A}[\![M]\!]_{\rho[x_i \mapsto r_i]_{i=1}^n, \iota, j} @ \text{index}(a \diamond \iota.j) \\
&\mid \langle j, \vec{r} \rangle \leftarrow \text{enum}([\vec{r} \mid [r_i \leftarrow \llbracket t_i \rrbracket]_{i=1}^n, \mathcal{N}[\![X]\!]_{\rho[x_i \mapsto r_i]_{i=1}^n})])])
\end{aligned}$$

As well as an environment, the current dynamic index is threaded through the semantics. Note that we use the ordinary semantics to evaluate values that cannot have any annotations (such as boolean tests in the last case).

It is straightforward to show by induction that the annotated semantics erases to the standard semantics:

THEOREM 19. *For any L we have $\text{erase}(\mathcal{A}[\![L]\!]) = \mathcal{N}[\![\text{erase}(L)]\!]$.*

D.2 Shredding and stitching annotated values

To define the semantics of shredded queries and packages, we use annotated values in which collections are annotated pairs of indexes and annotated values. Again, we allow annotations $@J$ on elements of collections to facilitate the proof of correctness. Typing rules for these values are shown in Figure 18.

Results	$s ::= [(I_1, w_1) @ J_1, \dots, (I_m, w_m) @ J_m]$
Shredded values	$w ::= c \mid r \mid I$
Rows	$r ::= \langle \ell_1 = w_1, \dots, \ell_n = w_n \rangle$
Indexes	I, J

Having defined suitably annotated versions of nested and shredded values, we now extend the shredding function to operate on nested values.

$$\begin{aligned}
\llbracket s \rrbracket_p &= \llbracket s \rrbracket_{\top}^* \diamond_{1,p} \\
\llbracket [v_i @ J_i]_{i=1}^n \rrbracket_{I,\epsilon}^* &= \langle I, \llbracket v_i \rrbracket_{J_i}^* @ J_i \rangle_{i=1}^n \\
\llbracket [v_i @ J_i]_{i=1}^n \rrbracket_{I,\iota.p}^* &= \text{concat}([\llbracket v_i \rrbracket_{J_i.p}^*]_{i=1}^n) \\
\llbracket \langle \ell_i = v_i \rangle_{i=1}^n \rrbracket_{I,\ell_i.p}^* &= \llbracket v_i \rrbracket_{I,p}^* \\
\llbracket c \rrbracket_I &= c \\
\llbracket \langle \ell_i = v_i \rangle_{i=1}^n \rrbracket_I &= \langle \ell_i = \llbracket v_i \rrbracket_I \rangle_{i=1}^n \\
\llbracket s \rrbracket_I &= I
\end{aligned}$$

Note that in the cases for bags, the index annotation J is passed as an argument to the recursive call: this is where we need the ghost indexes, to relate the semantics of nested and shredded queries.

We lift the nested result shredding function to build an (annotated) shredded value package in the same way that we did for

RESULT $\frac{[\vdash I_i : Index]_{i=1}^n \quad [\vdash w_i : F]_{i=1}^n \quad [\vdash J_i : Index]_{i=1}^n}{\vdash \langle I_1, w_1 @ J_1 \rangle, \dots, \langle I_n, w_n @ J_n \rangle : \text{Bag } \langle Index, F \rangle}$	CONSTANT $\frac{\Sigma(c) = O}{\vdash c : O}$	RECORD $\frac{[\vdash n_i : F_i]_{i=1}^n}{\vdash \langle \ell_i = w_i \rangle_{i=1}^n : \langle \ell_i = F_i \rangle_{i=1}^n}$
---	---	--

Figure 18: Typing rules for shredded values

nested types and nested queries.

$$\begin{aligned} \text{shred}_s(\text{Bag } A) &= \text{package}_{(\llbracket s \rrbracket _)}(\text{Bag } A) \\ \text{shred}_{s,p}(\text{Bag } A) &= \text{package}_{(\llbracket s \rrbracket _),p}(\text{Bag } A) \end{aligned}$$

We adjust the stitching function slightly to maintain annotations; the adjustment is consistent with its behaviour on unannotated values.

$$\begin{aligned} \text{stitch}(\hat{A}) &= \text{stitch}_{\top \circ 1}(\hat{A}) \\ \text{stitch}_c(O) &= c \\ \text{stitch}_r(\langle \ell_i : \hat{A}_i \rangle_{i=1}^n) &= \langle \ell_i = \text{stitch}_{r,cl_i}(\hat{A}_i) \rangle_{i=1}^n \\ \text{stitch}_I((\text{Bag } A)^s) &= [(\text{stitch}_w(\hat{A})) @ J \mid \langle I, w \rangle @ J \leftarrow s] \end{aligned}$$

The inner value parameter w to the auxiliary function $\text{stitch}_w(-)$ specifies which values to stitch along the current path.

We now show how to interpret shredded queries and query packages over shredded values. The semantics of shredded queries (including annotations) is given in Figure 17. The semantics of a shredded query package is a *shredded value package* containing indexed results for each shredded query. For each type A we define $\mathcal{H}[A] = \text{shred}_A(A)$ and for each flat-nested, closed $\vdash L : A$ we define $\mathcal{H}[L]_A : \mathcal{H}[A]$ as $\text{pmap}_{S[_]}(\text{shred}_L(A))$.

D.3 Main results

There are three key theorems. The first states that shredding commutes with the annotated semantics.

THEOREM 20. *If $\vdash L : \text{Bag } A$ then:*

$$\mathcal{H}[\llbracket L \rrbracket_{\text{Bag } A}] = \text{shred}_{\mathcal{A}[\llbracket L \rrbracket]}(\text{Bag } A)$$

In order to allow shredded results to be correctly stitched together, we need the indexes at the end of each path to a bag through a nested value to be unique. We define $\text{indexes}_p(v)$, the indexes of nested value v along path p as follows.

$$\begin{aligned} \text{indexes}_e(\langle v_i @ J_i \rangle_{i=1}^n) &= [J_i]_{i=1}^n \\ \text{indexes}_{\downarrow,p}(\langle v_i @ J_i \rangle_{i=1}^n) &= \text{concat}([\text{indexes}_p(v_i)]_{i=1}^n) \\ \text{indexes}_{\ell_i,p}(\langle \ell_i = v_i \rangle_{i=1}^n) &= \text{indexes}_p(v_i) \end{aligned}$$

We say that a nested value v is *well-indexed* (at type A) provided $\vdash v : A$, and for every path p in $\text{paths}(A)$, the elements of $\text{indexes}_p(v)$ are distinct.

LEMMA 21. *If $\vdash L : A$, then $\mathcal{A}[\llbracket L \rrbracket]$ is well-indexed at A .*

Our next theorem states that for well-indexed values, stitching is a left-inverse of shredding.

THEOREM 22. *If $\vdash s : \text{Bag } A$ and s is well-indexed at type $\text{Bag } A$ then $\text{stitch}(\text{shred}_s(\text{Bag } A)) = s$.*

Combining Theorem 20, Lemma 21 and Theorem 22 we obtain the main correctness result (see Figure 15).

THEOREM 23. *If $\vdash L : \text{Bag } A$ then: $\text{stitch}(\mathcal{H}[\llbracket L \rrbracket_{\text{Bag } A}]) = \mathcal{A}[\llbracket L \rrbracket]$, and in particular: $\text{erase}(\text{stitch}(\mathcal{H}[\llbracket L \rrbracket_{\text{Bag } A}])) = \mathcal{N}[\llbracket L \rrbracket]$.*

Not all possible indexing schemes satisfy Lemma 21. To identify those that do, we recall the function for computing the canonical indexes of a nested query.

$$\begin{aligned} \mathcal{I}[\llbracket L \rrbracket] &= \mathcal{I}[\llbracket L \rrbracket]_{\varepsilon,1} \\ \mathcal{I}[\llbracket \bigoplus_{i=1}^n C_i \rrbracket]_{\rho,\iota} &= \text{concat}([\mathcal{I}[C_i]_{\rho,\iota}]_{i=1}^n) \\ \mathcal{I}[\llbracket \langle \ell_i = M_i \rangle_{i=1}^n \rrbracket]_{\rho,\iota} &= \text{concat}([\mathcal{I}[M_i]_{\rho,\iota}]_{i=1}^n) \\ \mathcal{I}[\llbracket X \rrbracket]_{\rho,\iota} &= [] \\ \mathcal{I}[\llbracket \text{for } ([x_i \leftarrow t_i]_{i=1}^n \text{ where } X) \text{ return }^a M \rrbracket]_{\rho,\iota} &= \\ \text{concat}([a \diamond \iota.j :: \mathcal{I}[M]_{\rho[x_i \mapsto r_i]_{i=1}^n, \iota.j}} & \\ | \langle j, \vec{r} \rangle \leftarrow \text{enum}([\vec{r} \mid [r_i \leftarrow \llbracket t_i \rrbracket]_{i=1}^n, \mathcal{N}[X]_{\rho[x_i \mapsto r_i]_{i=1}^n}])]) & \end{aligned}$$

Note that $\mathcal{I}[_]$ resembles $\mathcal{A}[_]$, but instead of the nested value v it computes all indexes of v . An indexing function $\text{index} : Index \rightarrow A$ is *valid* with respect to the closed nested query L if it is injective and defined on every canonical index in $\mathcal{I}[\llbracket L \rrbracket]$.

LEMMA 24. *If index is valid for L then $\mathcal{A}[\llbracket L \rrbracket]$ is well-indexed.*

The only requirement on indexes in the proof of Theorem 4 is that nested values be well-indexed, hence the proof extends to any valid indexing scheme. The concrete, natural, and flat indexing schemes are all valid.

D.4 Detailed proofs

We first show that the inner shredding function $\llbracket _ \rrbracket$ commutes with the annotated semantics.

LEMMA 25. $\mathcal{S}[\llbracket \llbracket M \rrbracket_a \rrbracket]_{\rho,\iota} = \llbracket \mathcal{A}[M]_{\rho,\iota} \rrbracket_{a \diamond \iota}$

PROOF. By induction on the structure of M .

Case $x.\ell$:

$$\begin{aligned} &\mathcal{S}[\llbracket x.\ell \rrbracket_a]_{\rho,\iota} \\ &= \mathcal{S}[x.\ell]_{\rho,\iota} \\ &= \mathcal{A}[x.\ell]_{\rho,\iota} \\ &= \llbracket \mathcal{A}[x.\ell]_{\rho,\iota} \rrbracket_{a \diamond \iota} \end{aligned}$$

Case $c([X_i]_{i=1}^n)$:

$$\begin{aligned} &\mathcal{S}[\llbracket c([X_i]_{i=1}^n) \rrbracket_a]_{\rho,\iota} \\ &= \mathcal{S}[c(\llbracket [X_i]_{i=1}^n \rrbracket_a)]_{\rho,\iota} \\ &= \mathcal{A}[c(\llbracket [X_i]_{i=1}^n \rrbracket_a)]_{\rho,\iota} \\ &= \llbracket c \rrbracket([\mathcal{A}[\llbracket [X_i]_{i=1}^n \rrbracket_a]_{\rho,\iota}]_{i=1}^n) \\ &= \llbracket \mathcal{A}[c([X_i]_{i=1}^n)]_{\rho,\iota} \rrbracket_{a \diamond \iota} \end{aligned}$$

Case empty M :

$$\begin{aligned}
& \mathcal{S}[\llbracket \text{empty } M \rrbracket_a]_{\rho, \iota} \\
&= \mathcal{S}[\text{empty } \llbracket M \rrbracket_\epsilon]_{\rho, \iota} \\
&= \mathcal{A}[\text{empty } \llbracket M \rrbracket_\epsilon]_{\rho, \iota} \\
&= \llbracket \mathcal{A}[\text{empty } M] \rrbracket_{\rho, \iota} \llbracket a \diamond \iota \rrbracket
\end{aligned}$$

Case $\langle \ell = M \rangle$:

$$\begin{aligned}
& \mathcal{S}[\llbracket \langle \ell_i = M_i \rangle_{i=1}^n \rrbracket_a]_{\rho, \iota} \\
&= \mathcal{S}[\langle \ell_i = \llbracket M_i \rrbracket_a \rangle_{i=1}^n]_{\rho, \iota} \\
&= \langle \ell_i = \mathcal{S}[\llbracket M_i \rrbracket_a]_{\rho, \iota} \rangle_{i=1}^n \\
&= \text{(Induction hypothesis)} \\
&= \langle \ell_i = \llbracket \mathcal{A}[M_i] \rrbracket_{\rho, \iota} \rrbracket_{a \diamond \iota} \rangle_{i=1}^n \\
&= \llbracket \langle \ell_i = \mathcal{A}[M_i]_{\rho, \iota} \rangle_{i=1}^n \rrbracket_{a \diamond \iota} \\
&= \llbracket \mathcal{A}[\langle \ell_i = M_i \rangle_{i=1}^n]_{\rho, \iota} \rrbracket_{a \diamond \iota}
\end{aligned}$$

Case L :

$$\begin{aligned}
& \mathcal{S}[\llbracket L \rrbracket_a]_{\rho, \iota} \\
&= \mathcal{S}[a \diamond \text{in}]_{\rho, \iota} \\
&= a \diamond \iota \\
&= \llbracket \mathcal{A}[L]_{\rho, \iota} \rrbracket_{a \diamond \iota}
\end{aligned}$$

This completes the proof. \square

The first part of the following lemma allows us to run a shredded query by concatenating the results of running the shreds of its component comprehensions. Similarly, the second part allows us to shred the results of running a nested query by concatenating the shreds of the results of running its component comprehensions.

LEMMA 26.

1. $\mathcal{S}[\llbracket \biguplus_{i=1}^n C_i \rrbracket_{a,p}^*]_{\rho, \iota} = \text{concat}([\mathcal{S}[\llbracket C_i \rrbracket_{a,p}^*]_{\rho, \iota}]_{i=1}^n)$
2. $\llbracket \mathcal{A}[\biguplus_{i=1}^n C_i]_{\rho, \iota} \rrbracket_{a \diamond \iota, p}^* = \text{concat}([\llbracket \mathcal{A}[C_i]_{\rho, \iota} \rrbracket_{a \diamond \iota, p}^*]_{i=1}^n)$

PROOF.

1.

$$\begin{aligned}
& \mathcal{S}[\llbracket \biguplus_{i=1}^n C_i \rrbracket_{a,p}^*]_{\rho, \iota} \\
&= \mathcal{S}[\llbracket \biguplus_{i=1}^n \text{concat}([\llbracket C_i \rrbracket_{a,p}^*]_{i=1}^n) \rrbracket_{\rho, \iota} \\
&= \mathcal{S}[\llbracket \biguplus_{i=1}^n \text{concat}([\llbracket C_i \rrbracket_{a,p}^* \mid C \leftarrow \vec{C}]) \rrbracket_{\rho, \iota} \\
&= \text{concat}(\text{concat}([\mathcal{S}[\llbracket C' \rrbracket_{\rho, \iota} \mid C \leftarrow \vec{C}, C' \leftarrow \llbracket C \rrbracket_{a,p}^*]) \\
&= \text{concat}([\mathcal{S}[\llbracket \biguplus_{i=1}^n \llbracket C_i \rrbracket_{a,p}^* \mid C \leftarrow \vec{C} \rrbracket]_{\rho, \iota} \\
&= \text{concat}([\mathcal{S}[\llbracket \biguplus_{i=1}^n \llbracket C_i \rrbracket_{a,p}^* \rrbracket_{\rho, \iota}]]_{i=1}^n)
\end{aligned}$$

2.

$$\begin{aligned}
& \llbracket \mathcal{A}[\biguplus_{i=1}^n C_i]_{\rho, \iota} \rrbracket_{a \diamond \iota, p}^* \\
&= \llbracket \text{concat}([\mathcal{A}[C_i]_{\rho, \iota}]_{i=1}^n) \rrbracket_{a \diamond \iota, p}^* \\
&= \llbracket \text{concat}([\mathcal{A}[C]_{\rho, \iota} \mid C \leftarrow \vec{C}]) \rrbracket_{a \diamond \iota, p}^* \\
&= \text{concat}(\text{concat}([\llbracket s \rrbracket_{a \diamond \iota, p}^* \mid C \leftarrow \vec{C}, s \leftarrow \mathcal{A}[C]_{\rho, \iota}]) \\
&= \text{concat}([\llbracket \mathcal{A}[C]_{\rho, \iota} \rrbracket_{a \diamond \iota, p}^* \mid C \leftarrow \vec{C}]) \\
&= \text{concat}([\llbracket \mathcal{A}[C_i]_{\rho, \iota} \rrbracket_{a \diamond \iota, p}^*]_{i=1}^n)
\end{aligned}$$

This completes the proof. \square

We are now in a position to prove that the outer shredding function $\llbracket - \rrbracket$ commutes with the semantics.

LEMMA 27. $\mathcal{S}[\llbracket L \rrbracket_p]_{\rho, 1} = \llbracket \mathcal{A}[L]_{\rho, 1} \rrbracket_p$

PROOF. We prove the following:

1. $\mathcal{S}[\llbracket L \rrbracket_p]_{\rho, 1} = \llbracket \mathcal{A}[L]_{\rho, 1} \rrbracket_p$
2. $\mathcal{S}[\llbracket \biguplus C \rrbracket_{a,p}^*]_{\rho, \iota} = \llbracket \mathcal{A}[\biguplus C]_{\rho, \iota} \rrbracket_{a \diamond \iota, p}^*$
3. $\mathcal{S}[\llbracket M \rrbracket_{a,p}^*]_{\rho, \iota} = \llbracket \mathcal{A}[M]_{\rho, \iota} \rrbracket_{a \diamond \iota, p}^*$

The first equation is the result we require. Observe that it follows from (2) and Lemma 26. We now proceed to prove equations (2) and (3) by mutual induction on the structure of p .

There are only two cases for (2), as the $\ell_i.p$ case cannot apply.

Case ϵ :

$$\begin{aligned}
& \mathcal{S}[\llbracket \biguplus \text{for } (x \leftarrow \vec{t} \text{ where } X) \text{ return }^b M \rrbracket_{a,\epsilon}^*]_{\rho, \iota} \\
&= \text{(Definition of } \mathcal{S}[\llbracket - \rrbracket] \text{)} \\
&= \mathcal{S}[\text{for } (x \leftarrow \vec{t} \text{ where } X) \text{ return }^b \langle a \diamond \iota, \llbracket M \rrbracket_b \rangle]_{\rho, \iota} \\
&= \text{(Definition of } \mathcal{S}[\llbracket - \rrbracket] \text{)} \\
&= [\langle a \diamond \iota, \mathcal{S}[\llbracket M \rrbracket_b]_{\rho[\vec{x} \mapsto \vec{t}], \iota, i} \rangle @b \diamond \iota, i] \\
&= \text{(Lemma 25)} \\
&= [\langle a \diamond \iota, \llbracket \mathcal{A}[M]_{\rho[\vec{x} \mapsto \vec{t}], \iota, i} \rrbracket_{b \diamond \iota, i} \rangle @b \diamond \iota, i] \\
&= \text{(Definition of } \llbracket - \rrbracket^* \text{)} \\
&= \llbracket \mathcal{A}[\llbracket M \rrbracket_{\rho[\vec{x} \mapsto \vec{t}], \iota, i} @b \diamond \iota, i] \rrbracket_{a \diamond \iota, \epsilon}^* \\
&= \text{(Definition of } \mathcal{A}[\llbracket - \rrbracket] \text{)} \\
&= \llbracket \mathcal{A}[\text{for } (x \leftarrow \vec{t} \text{ where } X) \text{ return }^b M]_{\rho, \iota} \rrbracket_{a \diamond \iota, \epsilon}^*
\end{aligned}$$

Case $\downarrow.\epsilon$:

$$\begin{aligned}
& \mathcal{S}[\llbracket \text{for } (\overrightarrow{x \leftarrow t} \text{ where } X) \text{ return }^b M \rrbracket_{a, \downarrow, p}^*]_{\rho, \iota} \\
&= \text{(Definition of } \llbracket - \rrbracket^* \text{)} \\
& \mathcal{S}[\llbracket \text{for } (\overrightarrow{x \leftarrow t} \text{ where } X) C \mid C \leftarrow \llbracket M \rrbracket_{b, p}^* \rrbracket]_{\rho, \iota} \\
&= \text{(Definition of } \mathcal{S}[\llbracket - \rrbracket] \text{)} \\
& \text{concat}([\text{concat} \\
& \quad ([\mathcal{S}[C]_{\rho[\overrightarrow{x \mapsto b}], \iota, i} \\
& \quad \mid \langle i, \vec{v} \rangle \leftarrow \text{enum}([\vec{v} \mid v \leftarrow \llbracket t \rrbracket, \mathcal{S}[X]_{\rho[\overrightarrow{x \mapsto b}]}])]) \\
& \quad \mid C \leftarrow \llbracket M \rrbracket_{b, p}^*]) \\
&= \text{(Definition of } \mathcal{S}[\llbracket - \rrbracket] \text{)} \\
& \text{concat} \\
& \quad ([\mathcal{S}[\llbracket M \rrbracket_{b, p}^*]_{\rho[\overrightarrow{x \mapsto b}], \iota, i} \\
& \quad \mid \langle i, \vec{v} \rangle \leftarrow \text{enum}([\vec{v} \mid v \leftarrow \llbracket t \rrbracket, \mathcal{S}[X]_{\rho[\overrightarrow{x \mapsto b}]}])]) \\
&= \text{(Induction hypothesis (3))} \\
& \text{concat} \\
& \quad ([\llbracket \mathcal{A}[M] \rrbracket_{\rho[\overrightarrow{x \mapsto b}], \iota, i}^*]_{b, p} \\
& \quad \mid \langle i, \vec{v} \rangle \leftarrow \text{enum}([\vec{v} \mid v \leftarrow \llbracket t \rrbracket, \mathcal{S}[X]_{\rho[\overrightarrow{x \mapsto b}]}])]) \\
&= \text{(Definitions of } \mathcal{A}[\llbracket - \rrbracket \text{]} \text{ and } \llbracket - \rrbracket^* \text{)} \\
& \llbracket \mathcal{A}[\text{for } (\overrightarrow{x \leftarrow t} \text{ where } X) \text{ return }^b M] \rrbracket_{\rho, \iota}^*]_{a \diamond \iota, \downarrow, p}
\end{aligned}$$

There are three cases for (3).

Cases ϵ and $\downarrow.\epsilon$: follow from (2) by applying the two parts of Lemma 26 to the left and right-hand side respectively.

Case $\ell_i.\epsilon$:

$$\begin{aligned}
& \mathcal{S}[\llbracket \langle \ell = \vec{M} \rangle \rrbracket_{a, \ell_i, p}^*]_{\rho, \iota} \\
&= \text{(Definition of } \llbracket - \rrbracket^* \text{)} \\
& \mathcal{S}[\llbracket M_i \rrbracket_{a, p}^*]_{\rho, \iota} \\
&= \text{(Induction hypothesis (3))} \\
& \llbracket \mathcal{A}[M_i] \rrbracket_{\rho, \iota}^*]_{a \diamond \iota, p} \\
&= \text{(Definition of } \mathcal{A}[\llbracket - \rrbracket] \text{)} \\
& \llbracket \mathcal{A}[\langle \ell = \vec{M} \rangle] \rrbracket_{\rho, \iota}^*]_{a \diamond \iota, \ell_i, p}
\end{aligned}$$

This completes the proof. \square

We now lift Lemma 27 to shredded packages.

PROOF OF THEOREM 20. We need to show that if $\vdash L : \text{Bag } A$ then

$$\mathcal{H}[L]_A = \text{shred}_{\mathcal{A}[L]}(\text{Bag } A)$$

This is straightforward by induction on A , using Lemma 27. \square

We have proved that shredding commutes with the semantics. It remains to show that stitching after shredding is the identity on index-annotated nested results. We need two auxiliary notions: the *descendant* of a value at a path, and the *indexes* at the end of a path (which must be unique in order for stitching to work).

We define $\llbracket v \rrbracket_{p, w}$, the *descendant* of a value v at path p with respect to inner value w as follows.

$$\begin{aligned}
\llbracket v \rrbracket_{p, w} &= \llbracket v \rrbracket_{\top \diamond 1, p, w} \\
\llbracket v \rrbracket_{J, p, c} &= c \\
\llbracket v \rrbracket_{J, p, \langle \ell_i = w_i \rangle_{i=1}^n} &= \langle \ell_i = \llbracket v \rrbracket_{J, p, \ell_i, w_i} \rangle_{i=1}^n \\
\llbracket s \rrbracket_{J, \epsilon, I} &= \begin{cases} s, & \text{if } J = I \\ [], & \text{if } J \neq I \end{cases} \\
\llbracket [v_i @ J_i]_{i=1}^n \rrbracket_{J, \downarrow, p, I} &= \text{concat}([\llbracket v_i \rrbracket_{J_i, p, I} \rrbracket_{i=1}^n]) \\
\llbracket \langle \ell_i = v_i \rangle_{i=1}^n \rrbracket_{J, \ell_i, p, I} &= \llbracket v_i \rrbracket_{J, p, I}
\end{aligned}$$

Essentially, this extracts the part of v that corresponds to w . The inner value w allows us to specify a particular descendant as an index, or nested record of indexes; for uniformity it may also contain constants.

We can now formulate the following crucial technical lemma, which states that given the descendants of a result v at path $p.\downarrow$ and the shredded values of v at path p we can stitch them together to form the descendants at path p .

LEMMA 28 (KEY LEMMA). *If v is well-indexed and $\vdash \llbracket v \rrbracket_{J, p, I} : \text{Bag } A$, then*

$$\llbracket v \rrbracket_{J, p, I} = [\llbracket v \rrbracket_{J, p, \downarrow, w} @ I_{\text{in}} \mid \langle I_{\text{out}}, w \rangle @ I_{\text{in}} \leftarrow \llbracket v \rrbracket_{J, p}^*, I_{\text{out}} = I]$$

PROOF. First we strengthen the induction hypothesis to account for records. The generalised induction hypothesis is as follows.

If v is well-indexed and $\vdash \llbracket v \rrbracket_{J', p, \vec{\ell}, I} : \text{Bag } A$, then

$$\begin{aligned}
& [\llbracket v \rrbracket_{J', p, \downarrow, \vec{\ell}, w} @ I_{\text{in}} \mid \langle I_{\text{out}}, w \rangle @ I_{\text{in}} \leftarrow \llbracket v \rrbracket_{J', p, \vec{\ell}}^*, I_{\text{out}} = I] \\
&= [v'.\vec{\ell} @ I_{\text{in}} \mid v' @ I_{\text{in}} \leftarrow \llbracket v \rrbracket_{J', p, I}]
\end{aligned}$$

The proof now proceeds by induction on the structure of A and side-induction on the structure of p .

Case O :

Subcase ϵ : If $J' \neq I$ then both sides are empty lists. Suppose that $J' = I$.

$$\begin{aligned}
& [\llbracket s \rrbracket_{J', \downarrow, \vec{\ell}, c} @ I_{\text{in}} \\
& \quad \mid \langle I_{\text{out}}, c \rangle @ I_{\text{in}} \leftarrow \llbracket s \rrbracket_{J', \vec{\ell}}^*, I_{\text{out}} = I] \\
&= \text{(Definition of } \llbracket - \rrbracket \text{)} \\
& [c @ I_{\text{in}} \mid \langle I_{\text{out}}, c \rangle @ I_{\text{in}} \leftarrow \llbracket s \rrbracket_{J', \vec{\ell}}^*, I_{\text{out}} = I] \\
&= (J' = I) \\
& s \\
&= \text{(Definition of } \llbracket - \rrbracket \text{)} \\
& [v'.\vec{\ell} @ I_{\text{in}} \mid v' @ I_{\text{in}} \leftarrow \llbracket s \rrbracket_{J', \epsilon, I}]
\end{aligned}$$

Subcase $\ell_i.p$:

$$\begin{aligned}
& [\llbracket \langle \ell = \vec{v} \rangle \rrbracket_{J', \ell_i, p, \downarrow, \vec{\ell}, c} @ I_{\text{in}} \\
& \quad \mid \langle I_{\text{out}}, c @ I_{\text{in}} \rangle \leftarrow \llbracket \langle \ell = \vec{v} \rangle \rrbracket_{J', \ell_i, p, \vec{\ell}}^*, I_{\text{out}} = I] \\
&= \text{(Definition of } \llbracket - \rrbracket \text{)} \\
& [c @ I_{\text{in}} \mid \langle I_{\text{out}}, c @ I_{\text{in}} \rangle \leftarrow \llbracket \langle \ell = \vec{v} \rangle \rrbracket_{J', \ell_i, p, \vec{\ell}}^*, I_{\text{out}} = I] \\
&= \text{(Definition of } \llbracket - \rrbracket^* \text{)} \\
& [c @ I_{\text{in}} \mid \langle I_{\text{out}}, c @ I_{\text{in}} \rangle \leftarrow \llbracket v_i \rrbracket_{J', p, \vec{\ell}}^*, I_{\text{out}} = I] \\
&= \text{(Definition of } \llbracket - \rrbracket \text{)} \\
& [\llbracket v_i \rrbracket_{J', p, \downarrow, \vec{\ell}, c} @ I_{\text{in}} \mid \langle I_{\text{out}}, c \rangle @ I_{\text{in}} \leftarrow \llbracket v_i \rrbracket_{J', p, \vec{\ell}}^*, I_{\text{out}} = I] \\
&= \text{(Induction hypothesis)} \\
& [v'.\vec{\ell} @ I_{\text{in}} \mid v' @ I_{\text{in}} \leftarrow \llbracket v_i \rrbracket_{J', p, I}] \\
&= \text{(Definition of } \llbracket - \rrbracket^* \text{)} \\
& [v'.\vec{\ell} @ I_{\text{in}} \mid v' @ I_{\text{in}} \leftarrow \llbracket \langle \ell = \vec{v} \rangle \rrbracket_{J', \ell_i, p, I}]
\end{aligned}$$

Subcase $\downarrow.p$:

$$\begin{aligned}
& [\llbracket [v_i @ J_i]_{i=1}^n \rrbracket_{J', \downarrow, p, \downarrow, \vec{\ell}, c} @ I_{\text{in}} \\
& \quad \mid \langle I_{\text{out}}, c \rangle @ I_{\text{in}} \leftarrow \llbracket [v_i @ J_i]_{i=1}^n \rrbracket_{J', \downarrow, p, \vec{\ell}}^*, I_{\text{out}} = I] \\
&= \text{(Definitions of } \llbracket - \rrbracket \text{ and } \llbracket - \rrbracket^* \text{)} \\
& \text{concat}([c @ I_{\text{in}} \mid \langle I_{\text{out}}, c @ I_{\text{in}} \rangle \leftarrow \llbracket v_i \rrbracket_{J_i, p, \vec{\ell}}^*, I_{\text{out}} = I]_{i=1}^n) \\
&= \text{(Induction hypothesis)} \\
& \text{concat}([v'.\vec{\ell} @ I_{\text{in}} \mid v' @ I_{\text{in}} \leftarrow \llbracket v_i \rrbracket_{J_i, p} \rrbracket_{i=1}^n]) \\
&= \text{(Definition of } \llbracket - \rrbracket \text{)} \\
& [v'.\vec{\ell} @ I_{\text{in}} \mid v' @ I_{\text{in}} \leftarrow \llbracket [v @ J]_{i=1}^n \rrbracket_{J', p}]
\end{aligned}$$

Case $\langle \rangle$: The proof is the same as for base types with the constant c replaced by $\langle \rangle$.

Case $\langle \ell : \vec{A} \rangle$ where $|\vec{\ell}| \geq q$: We rely on the functions $zip_{\vec{\ell}}$, for transforming a record of lists of equal length to a list of records, and $unzip_{\vec{\ell}}$, for transforming a list of records to a record of lists of equal length. In fact we require special versions of zip and $unzip$ that handle annotations, such that zip takes a record of lists of equal length whose annotations must be in sync, and $unzip$ returns such a record.

$$\begin{aligned} zip_{\vec{\ell}}(\ell_i = []_{i=1}^n) &= [] \\ zip_{\vec{\ell}}(\ell_i = v_i @ J_i :: s_i)_{i=1}^n &= \langle \ell_i = v_i @ J_i :: zip_{\vec{\ell}}(\ell_i = s_i)_{i=1}^n \rangle \\ unzip_{\vec{\ell}}(s) &= \langle \ell_i = [v_i @ J_i \mid v @ J_i \leftarrow s]_{i=1}^n \rangle \end{aligned}$$

If $\vec{\ell}$ is a non-empty list of column labels then $zip_{\vec{\ell}}$ is the inverse of $unzip_{\vec{\ell}}$.

$$zip_{\vec{\ell}}(unzip_{\vec{\ell}}(s)) = s, \quad \text{if } |\vec{\ell}| \geq 1$$

$$\begin{aligned} & \llbracket v \ll_{J', p, \downarrow, \vec{\ell}, w} @ I_{in} \mid \langle I_{out}, w \rangle @ I_{in} \leftarrow \llbracket v \rrbracket_{J', p, \vec{\ell}}^*, I_{out} = I \rrbracket \\ &= \text{(Definition of } \ll - \gg \text{)} \\ & \llbracket \langle \ell_i = \llbracket v \rrbracket_{J', p, \downarrow, \vec{\ell}, w} @ I_{in} \mid \langle I_{out}, w @ I_{in} \rangle \leftarrow \llbracket v \rrbracket_{J', p, \vec{\ell}}^*, I_{out} = I \rrbracket \\ &= \text{(Definition of } zip \text{)} \\ & zip_{\vec{\ell}}(\ell_i = [v_{\ell_i, w} @ I_{in} \mid \langle I_{out}, w @ I_{in} \rangle \leftarrow \llbracket v \rrbracket_{J', p, \vec{\ell}}^*, I_{out} = I]_{i=1}^n) \\ & \quad \text{where } v_{\ell_i, w} \text{ stands for } \llbracket v \rrbracket_{J', p, \downarrow, \vec{\ell}, \ell_i, w} \\ &= \text{(Induction hypothesis)} \\ & zip_{\vec{\ell}}(\ell_i = [v'.\vec{\ell}'.\ell_i @ I_{in} \mid v' @ I_{in} \leftarrow \llbracket v \rrbracket_{J', p, I}^*]_{i=1}^n) \\ &= \text{(Definition of } zip \text{)} \\ & [v'.\vec{\ell}' @ I_{in} \mid v' @ I_{in} \leftarrow \llbracket v \rrbracket_{J', p, I}^*] \end{aligned}$$

Case Bag \vec{A} : Subcase ϵ : If $J' \neq I$ then both sides of the equation are equivalent to the empty bag. Suppose $J' = I$.

$$\begin{aligned} & \llbracket [v_i @ J_i]_{i=1}^n \ll_{J', \downarrow, \vec{\ell}, I_{in}} @ I_{in} \mid \langle I_{out}, I_{in} \rangle @ I_{in} \leftarrow \llbracket [v_i @ J_i]_{i=1}^n \rrbracket_{J', \vec{\ell}}^*, I_{out} = I \rrbracket \\ &= \text{(Definition of } \ll - \gg \text{)} \\ & \llbracket concat([v_i \ll_{J_i, \vec{\ell}, I_{in}} @ I_{in}]_{i=1}^n) \mid \langle I_{out}, I_{in} \rangle @ I_{in} \leftarrow \llbracket [v_i @ J_i]_{i=1}^n \rrbracket_{J', \vec{\ell}}^*, I_{out} = I \rrbracket \\ &= \llbracket \llbracket [v_i @ J_i]_{i=1}^n \rrbracket_{J', \vec{\ell}}^* = \langle J', J_1 \rangle @ J_1, \dots, \langle J', J_n \rangle @ J_n \rrbracket \\ & \llbracket concat([v_i \ll_{J_i, \vec{\ell}, I_{in}} @ I_{in}]_{i=1}^n) @ I_{in} \mid I_{in} \leftarrow \vec{J}, J' = I \rrbracket \\ &= \text{(} J' = I \text{)} \\ & \llbracket concat([v_i \ll_{J_i, \vec{\ell}, I_{in}} @ I_{in}]_{i=1}^n) @ I_{in} \mid I_{in} \leftarrow \vec{J} \rrbracket \\ &= \text{(Definition of } \ll - \gg \text{)} \\ & \llbracket concat([v_i \ll_{J_i, \epsilon, I_{in}} @ I_{in} \mid v_i @ J_i \leftarrow [v_i @ J_i]_{i=1}^n] @ I_{in} \mid I_{in} \leftarrow \vec{J} \rrbracket \\ &= \text{(Definition of } \ll - \gg \text{)} \\ & \llbracket concat([v_i.\vec{\ell}' @ I_{in} \mid v_i @ J_i \leftarrow [v_i @ J_i]_{i=1}^n, J_i = I_{in}] @ I_{in} \mid I_{in} \leftarrow \vec{J} \rrbracket \\ &= \text{(} v \text{ is well-indexed)} \\ & [v_i.\vec{\ell}' @ I_{in} \mid v_i @ I_{in} \leftarrow [v_i @ J_i]_{i=1}^n] \\ &= \text{(Definition of } \ll - \gg \text{ and } J' = I \text{)} \\ & [v_i.\vec{\ell}' @ I_{in} \mid v_i @ I_{in} \leftarrow \llbracket [v_i @ J_i]_{i=1}^n \rrbracket_{J', \epsilon, I} \rrbracket \end{aligned}$$

Subcase $\ell_i.p$:

$$\begin{aligned} & \llbracket \langle \ell_i = v_i \rangle_{i=1}^n \ll_{J', \ell_i.p, \downarrow, \vec{\ell}', I_{in}} @ I_{in} \mid \langle I_{out}, I_{in} \rangle @ I_{in} \leftarrow \llbracket \langle \ell_i = v_i \rangle_{i=1}^n \rrbracket_{J', \ell_i.p, \vec{\ell}'}^*, I_{out} = I \rrbracket \\ &= \text{(Definitions of } \ll - \gg \text{ and } \llbracket - \rrbracket^* \text{)} \\ & \llbracket v_i \ll_{J', p, \downarrow, \vec{\ell}', I_{in}} @ I_{in} \mid \langle I_{out}, I_{in} \rangle @ I_{in} \leftarrow \llbracket v_i \rrbracket_{J', p, \vec{\ell}'}^*, I_{out} = I \rrbracket \\ &= \text{(Induction hypothesis)} \\ & [v'.\vec{\ell}' @ I_{in} \mid v' @ I_{in} \leftarrow \llbracket v_i \rrbracket_{J', p, I}^*] \\ &= \text{(Definition of } \ll - \gg \text{)} \\ & [v'.\vec{\ell}' @ I_{in} \mid v' @ I_{in} \leftarrow \llbracket \langle \ell_i = v_i \rangle_{i=1}^n \rrbracket_{J', \ell_i.p, I} \rrbracket \end{aligned}$$

Subcase $\downarrow.p$:

$$\begin{aligned} & \llbracket [v_i @ J_i]_{i=1}^n \ll_{J', \downarrow.p, \downarrow, \vec{\ell}, I_{in}} @ I_{in} \mid \langle I_{out}, I_{in} \rangle @ I_{in} \leftarrow \llbracket [v_i @ J_i]_{i=1}^n \rrbracket_{I_{out}, \downarrow.p, \vec{\ell}}^*, I_{out} = I \rrbracket \\ &= \text{(Definitions of } \ll - \gg \text{ and } \llbracket - \rrbracket^* \text{)} \\ & \llbracket concat([v_i \ll_{J_i, p, \downarrow, \vec{\ell}, I_{in}} @ I_{in}]_{i=1}^n) \mid \langle I_{out}, I_{in} \rangle @ I_{in} \leftarrow concat([v_i \rrbracket_{J_i, p, \vec{\ell}}^*]_{i=1}^n), I_{out} = I \rrbracket \\ &= \text{(Expanding } concat([v_i \rrbracket_{J_i, p, \vec{\ell}}^*]_{i=1}^n) \text{)} \\ & \llbracket concat([v_i \ll_{J_i, p, \downarrow, \vec{\ell}, I_{in}} @ I_{in}]_{i=1}^n) \mid v_i @ J_i \leftarrow [v_i @ J_i]_{i=1}^n, \langle I_{out}, I_{in} \rangle @ I_{in} \leftarrow \llbracket v_i \rrbracket_{J_i, p, \vec{\ell}}^*, I_{out} = I \rrbracket \\ &= \text{(} v \text{ is well-indexed)} \\ & \llbracket v_i \ll_{J_i, p, \downarrow, \vec{\ell}, I_{in}} @ I_{in} \mid v_i @ J_i \leftarrow [v_i @ J_i]_{i=1}^n, \langle I_{out}, I_{in} \rangle @ I_{in} \leftarrow \llbracket v_i \rrbracket_{J_i, p, \vec{\ell}}^*, I_{out} = I \rrbracket \\ &= \text{(Comprehension } \rightarrow \text{ concatenation)} \\ & concat \\ & \llbracket [v_i \ll_{J_i, p, \downarrow, \vec{\ell}, I_{in}} @ I_{in} \mid \langle I_{out}, I_{in} \rangle @ I_{in} \leftarrow \llbracket v_i \rrbracket_{J_i, p, \vec{\ell}}^*, I_{out} = I]_{i=1}^n \rrbracket \\ &= \text{(Induction hypothesis)} \\ & concat([v'.\vec{\ell}' @ I_{in} \mid v' @ I_{in} \leftarrow \llbracket v_i \rrbracket_{J_i, \downarrow.p, I}^*]_{i=1}^n) \\ &= \text{(Concatenation } \rightarrow \text{ comprehension)} \\ & [v'.\vec{\ell}' @ I_{in} \mid v' @ I_{in} \leftarrow \llbracket [v_i @ J_i]_{i=1}^n \rrbracket_{J', \downarrow.p, I} \rrbracket \end{aligned}$$

This completes the proof. \square

The proof of this lemma is the only part of the formalisation that makes use of values being well-indexed. Stitching shredded results together does not depend on any other property of indexes, thus any representation of indexes that yields unique indexes suffices.

THEOREM 29. *If s is well-indexed and $\vdash \llbracket s \rrbracket_{p, w} : A$ then $stitch_w(shred_{s, p}(A)) = \llbracket s \rrbracket_{p, w}$.*

PROOF. By induction on the structure of A .

Case O :

$$\begin{aligned} & stitch_c(shred_{s, p}(O)) \\ &= \\ & c \\ &= \\ & \llbracket s \rrbracket_{p, c} \end{aligned}$$

Case $\langle \ell = \vec{A} \rangle$:

$$\begin{aligned}
& \text{stitch}_{\langle \ell_i = w_i \rangle_{i=1}^n} (\text{shred}_{s,p}(\langle \ell_i : A \rangle_{i=1}^n)) \\
= & \\
& \langle \ell_i = \text{stitch}_{w_i}(\text{shred}_{s,p,l}(A)) \rangle_{i=1}^n \\
= & \quad (\text{Induction hypothesis}) \\
& \langle \ell_i = \llbracket s \rrbracket_{p,l,w_i} \rangle_{i=1}^n \\
= & \\
& \llbracket s \rrbracket_{p, \langle \ell_i = w_i \rangle_{i=1}^n}
\end{aligned}$$

Case Bag A:

$$\begin{aligned}
& \text{stitch}_I(\text{shred}_{s,p}(\text{Bag } A)) \\
= & \\
& \text{stitch}_I(((\text{shred}_{s,p,\downarrow}(A))^{\llbracket s \rrbracket_p})) \\
= & \\
& [(\text{stitch}_n(\text{shred}_{s,p,\downarrow}(A)) @ I_{\text{in}}) \\
& \quad | \langle I_{\text{out}}, w \rangle @ I_{\text{in}} \leftarrow \llbracket s \rrbracket_p^*, I_{\text{out}} = I] \\
= & \quad (\text{Induction hypothesis}) \\
& [\llbracket s \rrbracket_{p,\downarrow,w} @ I_{\text{in}} | \langle I_{\text{out}}, w \rangle @ I_{\text{in}} \leftarrow \llbracket s \rrbracket_p^*, I_{\text{out}} = I] \\
= & \quad (\text{Lemma 28}) \\
& \llbracket s \rrbracket_{p,I}
\end{aligned}$$

This completes the proof. \square

PROOF OF THEOREM 22. By Theorem 29, setting $p = \epsilon$ and $w = \top \diamond 1$. \square

We now obtain the main result.

PROOF OF THEOREM 4. Recall the statement of Theorem 4: we need to show that $\vdash L : \text{Bag } A$ then:

$$\text{stitch}(\mathcal{H}[\llbracket L \rrbracket]_{\text{Bag } A}) = \text{stitch}(\text{shred}_{\mathcal{A}[\llbracket L \rrbracket]}(\text{Bag } A)) = \mathcal{A}[\llbracket L \rrbracket]$$

The first equation follows immediately from Theorem 20, and the second from Lemma 21 and Theorem 22.

Furthermore, applying *erase* to both sides we have:

$$\text{erase}(\text{stitch}(\mathcal{H}[\llbracket L \rrbracket]_{\text{Bag } A})) = \text{erase}(\mathcal{A}[\llbracket L \rrbracket]) = \mathcal{N}[\llbracket L \rrbracket]$$

where the second step follows by Theorem 19. \square

E. RECORD FLATTENING

Flat types. For simplicity we extend base types to include the unit type $\langle \rangle$. This allows us to define an entirely syntax-directed *unflattening* translation from flat record values to nested record values.

$$\begin{array}{ll}
\text{Types} & A, B ::= \text{Bag } \langle \vec{\ell} : \vec{O} \rangle \\
\text{Base types} & O ::= \text{Int} \mid \text{Bool} \mid \text{String} \mid \langle \rangle
\end{array}$$

The Links implementation diverges slightly from the presentation here. Rather than treating the unit type as a base type, it relies on type information to construct units when unflattening values.

Flat terms.

$$\begin{array}{ll}
\text{Query terms} & L, M ::= \biguplus \vec{C} \\
\text{Comprehensions} & C ::= \text{let } q = S \text{ in } S' \\
\text{Subqueries} & S ::= \text{for } (\vec{G} \text{ where } X) \text{ return } R \\
\text{Data sources} & u ::= t \mid q \\
\text{Generators} & G ::= x \leftarrow u \\
\text{Inner terms} & N ::= X \mid \text{index} \\
\text{Record terms} & R ::= \langle \vec{\ell} = \vec{N} \rangle \\
\text{Base terms} & X ::= x.\ell \mid c(\vec{X}) \mid \text{empty } L
\end{array}$$

Flattening types. The record flattening function $(-)^{\succ}$ flattens record types.

$$\begin{aligned}
(\text{Bag } A)^{\succ} &= \text{Bag } (A^{\succ}) \\
O^{\succ} &= \langle \bullet : O \rangle \\
\langle \rangle^{\succ} &= \langle \bullet : \langle \rangle \rangle \\
\langle \vec{\ell} : \vec{F} \rangle^{\succ} &= \langle [(\ell_i - \ell') : O \mid i \leftarrow \text{dom}(\vec{\ell}), (\ell' : O) \leftarrow F_i]^{\succ} \rangle
\end{aligned}$$

Labels in nested records are concatenated with the labels of their ancestors. Base and unit types are lifted to 1-ary records (with a special \bullet field) for uniformity and to aid with reconstruction of nested values from flattened values.

Flattening terms. The record flattening function $(-)^{\succ}$ is defined on terms as well as types.

$$\begin{aligned}
(\biguplus_{i=1}^n C_i)^{\succ} &= \biguplus_{i=1}^n (C_i)^{\succ} \\
(\text{let } q = S_{\text{out}} \text{ in } S_{\text{in}})^{\succ} &= \text{let } q = S_{\text{out}}^{\succ} \text{ in } S_{\text{in}}^{\succ} \\
(\text{for } (\vec{G} \text{ where } X) \text{ return } N)^{\succ} &= \text{for } (\vec{G} \text{ where } X^{\dagger}) \text{ return } N^{\succ}
\end{aligned}$$

$$\begin{aligned}
X^{\succ} &= \langle \bullet = X^{\dagger} \rangle \\
\langle \rangle^{\succ} &= \langle \bullet = \langle \rangle \rangle \\
(\langle \ell_i = N_i \rangle_{i=1}^m)^{\succ} &= \langle \ell_i - \ell'_j = X_j \rangle_{(i=1, j=1)}^{(m, n_i)}, \\
&\quad \text{where } N_i^{\succ} = \langle \ell'_j = X_j \rangle_{j=1}^{n_i}
\end{aligned}$$

$$\begin{aligned}
(x.\ell_1 \dots \ell_n)^{\dagger} &= x.\ell_1 \sqcup \dots \sqcup \ell_n \bullet \\
c(X_1, \dots, X_n)^{\dagger} &= c((X_1)^{\dagger}, \dots, (X_n)^{\dagger}) \\
(\text{empty } L)^{\dagger} &= \text{empty } L^{\succ}
\end{aligned}$$

The auxiliary $(-)^{\dagger}$ function flattens n -ary projections.

Type soundness.

$$\vdash L : A \Rightarrow \vdash L^{\succ} : A^{\succ}$$

Unflattening record values.

$$[r_1, \dots, r_n]^{\prec} = [(r_1)^{\prec}, \dots, (r_n)^{\prec}]$$

$$\begin{aligned} \langle \bullet = c \rangle^{\prec} &= c \\ \langle \bullet = \langle \rangle \rangle^{\prec} &= \langle \rangle \\ (\langle \ell_i - \ell'_j = c_j \rangle_{(i=1, j=1)}^{(m, n_i)})^{\prec} &= \langle \ell_i = (\langle \ell'_j = c_j \rangle_{j=1}^{n_i})^{\prec} \rangle_{i=1}^m \end{aligned}$$

Type soundness.

$$\vdash L : A \quad \Rightarrow \quad \vdash L^{\prec} : A^{\prec}$$

Correctness

PROPOSITION 30. *If L is a let-inserted query and $\vdash L : \text{Bag } A$, then*

$$(\mathcal{L}[\llbracket L^{\succ} \rrbracket])^{\prec} = \mathcal{L}[\llbracket L \rrbracket]$$